



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Riccius, L., Rocha, I. B. C. M., Bierkens, J., Kekkonen, H., & van der Meer, F. P. (2026). Integration of active learning and MCMC sampling for efficient Bayesian calibration of mechanical properties. *European Journal of Mechanics, A/Solids*, 117, Article 106015. <https://doi.org/10.1016/j.euromechsol.2026.106015>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.



Full length article

Integration of active learning and MCMC sampling for efficient Bayesian calibration of mechanical properties

Leon Riccius ^a*, Iuri B.C.M. Rocha ^a, Joris Bierkens ^b, Hanne Kekkonen ^b, Frans P. van der Meer ^a^a Delft University of Technology, Faculty of Civil Engineering and Geosciences, Delft, The Netherlands^b Delft University of Technology, Delft Institute of Applied Mathematics, Delft, The Netherlands

ARTICLE INFO

Keywords:

Bayesian inference
 Markov chain Monte Carlo
 Surrogate modelling
 Gaussian process regression
 Computational mechanics

ABSTRACT

Recent advancements in Markov chain Monte Carlo (MCMC) sampling and surrogate modelling have significantly enhanced the feasibility of Bayesian analysis across engineering fields. However, the selection and integration of surrogate models and cutting-edge MCMC algorithms, often depend on ad-hoc decisions. A systematic assessment of their combined influence on accuracy and efficiency is notably lacking. The present work offers a comprehensive comparative study, employing a scalable case study in computational mechanics focused on the inference of spatially varying material parameters, that sheds light on the impact of methodological choices for surrogate modelling and sampling. We show that a priori training of the surrogate model introduces large errors in the posterior estimation even in low to moderate dimensions. We introduce a simple active learning strategy based on the path of the MCMC algorithm that is superior to all a priori trained models, and determine its training data requirements. We demonstrate that the choice of the MCMC algorithm has only a small influence on the amount of training data but no significant influence on the accuracy of the resulting surrogate model. Further, we show that the accuracy of the posterior estimation largely depends on the surrogate model, but not even a tailored surrogate guarantees convergence of the MCMC. Finally, we identify the forward model as the bottleneck in the inference process, not the MCMC algorithm. While related works focus on employing advanced MCMC algorithms, we demonstrate that the training data requirements render the surrogate modelling approach infeasible before the benefits of these gradient-based MCMC algorithms on cheap models can be reaped.

1. Introduction

In engineering mechanics, computational models, such as those based on the finite element method (FEM), are used to simulate the mechanical response of structures and, where required, the underlying material behaviour. The number of parameters for these models can be large, for instance if the material properties vary in space. These parameters often represent an internal state of the model that is difficult or impossible to measure directly. Therefore, an inverse problem must be solved to estimate these model parameters from indirect experimental measurements, commonly referred to as observations. In the nonelastic regime, these inverse problems are typically non-linear, non-convex, and ill-posed, but they can be regularised elegantly by adopting a Bayesian approach. This approach to inverse modelling employs Bayes' theorem to update the distribution of the parameters based on the observed data (Dobrilla et al., 2023; Marsili et al., 2017). On the downside, Bayesian inverse modelling is computationally expensive, preventing its widespread application to large-scale problems. A good

overview of general constitutive model calibration methods is given in Römer et al. (2025), and for Bayesian methods specifically in Kiran et al. (2025).

Recent advancements in Markov chain Monte Carlo (MCMC) sampling and surrogate modelling have substantially improved the applicability of Bayesian analysis in these scenarios. Despite the potential of these innovations, the selection of surrogate models and integration with state-of-the-art MCMC algorithms often rely on ad-hoc choices. A systematic evaluation of their combined impact on overall performance remains scarce. This lack of clarity obscures the understanding of how each component impacts the overall effectiveness of the inference process. Addressing this gap, our study conducts a comparative analysis to elucidate the effects of the methodological decisions regarding surrogate modelling and sampling.

In Bayesian inference, a prior distribution for the parameters is assumed and combined with a likelihood function. The likelihood measures the agreement of the model predictions with the observed system

* Corresponding author.

E-mail address: l.f.riccius@tudelft.nl (L. Riccius).

response. Following the Bayesian formalism, a posterior distribution for the parameters is obtained. For instance, when inferring material properties from measurements, the posterior distribution must be approximated due to the non-linear nature of the likelihood function, e.g. via sampling from the joint distribution of data and parameters. The Bayesian analysis via MCMC entails the following computational steps: Generating a new state, evaluating the forward model, computing the likelihood, and accepting or rejecting the proposal. The bottleneck of the sampling effort is the evaluation of the forward model, which, in case of FEM-based inverse modelling, is an expensive simulator that maps the parameters to the observables. The state generation and acceptance steps are implemented by the sampling algorithm and determine how often the model and likelihood must be evaluated. The most commonly used sampling algorithm is the random walk Metropolis (RWM) algorithm due to its simplicity and guaranteed asymptotic convergence (Mengersen and Tweedie, 1996). However, its sampling efficiency decreases with the dimensionality of the parameter space (Roberts and Rosenthal, 2001).

One way to reduce the number of costly forward passes is by employing sampling algorithms that more efficiently produce uncorrelated samples from the posterior distribution. The first-order Metropolis adjusted Langevin algorithm (Roberts and Stramer, 2002), Hamiltonian Monte Carlo (Duane et al., 1987), or their second-order manifold variants (Girolami and Calderhead, 2011), as well as piecewise deterministic Markov process based algorithms (Bierkens et al., 2019) maintain high sampling efficiency in high-dimensional settings. Empirical studies of their performance on engineering problems (Girolami and Calderhead, 2011; Chong and Lam, 2017; Goodman and Weare, 2010) suggests their theoretical performance in terms of convergence rates on canonical distributions (Gelman et al., 1997; Roberts and Rosenthal, 1998) translates well to practical settings. However, these algorithms require the gradient of the log-joint density, or even higher derivatives thereof, quantities which are at best difficult to obtain for large scale FEM models. Ensemble-based methods (Ching and Chen, 2007; Lye et al., 2022; Straub and Papaioannou, 2015) perform well on certain classes of problems but still require at least $\mathcal{O}(10^4)$ model evaluations to ensure convergent posterior statistics.

The other increasingly common strategy to reduce the computational cost of the sampling effort is the application of surrogate models. Surrogate models exploit the redundancy in the model predictions required for collecting posterior samples. Instead of running the expensive simulator in each MCMC step, a cheap-to-evaluate surrogate model is trained on a small number of solver calls.

Such a surrogate model, is constructed in Hu et al. (2024), Thomas et al. (2022), Deveney et al. (2023). In Wu et al. (2020), e.g., a neural network surrogate is trained to replace an iterative mean-field homogenisation algorithm for material property inference in the presence of plastic deformations. All these approaches have in common that multiple models must be trained if the dataset consists of observations originating from fundamentally different conditions. The three distinct test scenarios in Wu et al. (2020) motivate the training of three neural networks to reduce the complexity of individual tasks. Further, the works employing polynomial chaos and Gaussian process (GP) models do not consider correlations between the outputs and train one model per scalar observation. While model training is typically not the computational bottleneck in the settings we consider here, tweaking the hyperparameters for the individual models can be cumbersome and time-consuming for the analyst.

After the model is evaluated, its responses are funnelled into the likelihood function to produce a scalar measure for the agreements of observations and model predictions for a given parameter vector. This function can then be approximated by a single surrogate model (Chen et al., 2023; Drovandi et al., 2018; del Val et al., 2022). However, targeting the likelihood function also introduces certain complexities. The likelihood incorporates all model non-linearity at once and might be more difficult of a target than the individual model predictions.

Moreover, the likelihood has its own parameters, e.g. the observation noise, that one might be uncertain about. These parameters of the observation model can be inferred alongside the model parameters, but this increases the dimensionality of the input space for the surrogate model. Despite these challenges, we prioritise the likelihood approach in our study. This decision is driven by its broader applicability across various domains, as it allows for a more generalised framework for modelling in contrast to methods tailored to specific physical models. It should be noted that multi-fidelity MCMC (Zhang et al., 2018) and non-linear dimensionality reduction techniques (Dasgupta et al., 2024) have shown promising results for Bayesian inverse modelling. However, these methods come with their own set of modelling challenges and therefore remain beyond the scope of this work.

For any choice of surrogate model, training data selection is crucial for the accuracy of the approximation. Most commonly, an a priori strategy is employed. The training points are either sampled from the prior, generated by Latin hypercube sampling (Hu et al., 2024; Thomas et al., 2022; Wu et al., 2020; Chen et al., 2023), or spread out on a grid (del Val et al., 2022). These a priori strategies suffer from the curse of dimensionality, i.e. their performance degrades rapidly with increasing dimensionality of the parameter space.

The model accuracy is typically measured by the mean squared error (MSE) on a validation set. This set is generated with the same strategy as the training data. However, this MSE is not necessarily a good indicator of the quality of the surrogate model for a Bayesian analysis — the surrogate model is only required to be accurate in regions of high posterior density. Active learning strategies, such as Deveney et al. (2023), Kandasamy et al. (2015), Drovandi et al. (2018), Dinkel et al. (2024), probe the posterior distribution to collect more informative data points. Further, surrogate models are generally not only cheap to evaluate, but also cheap to differentiate, enabling the application of gradient-based sampling algorithms.

The aforementioned works which focus on surrogate models for Bayesian inference identify MCMC sampling as a bottleneck and employ a variety of algorithms: RWM in Wu et al. (2020), del Val et al. (2022); HMC or its No U-Turn (NUTS) variant (Deveney et al., 2023; Hu et al., 2024; Thomas et al., 2022); TMCMC (Chen et al., 2023). However, the specific choices are rarely justified or even commented on. A good overview of the *mix and match* regarding surrogate and sampling algorithm can be found in the review of Hou et al. (2021) on Bayesian inference for building energy models. The review contains an assessment both of surrogate models and sampling algorithms in isolation, but does not cover the integration of both components into a single framework. Especially for active learning strategies, the interplay between the surrogate model and the sampling algorithm is crucial: the sampler determines where the surrogate model is queried, and the surrogate models in turn affects the acceptance probability, and the subsequent proposal generation.

In this work, we investigate the interplay of surrogate modelling and sampling algorithms for Bayesian inference. We focus on the inference of material parameters in the non-linear regime of a computational mechanics model. We introduce a scalable case study to create a series of inference problems with comparable complexity but increasing stochastic dimensionality. We employ a GP surrogate model for the likelihood and propose to follow the MCMC chain to collect training data. This form of active learning is based on uncertainty of the model prediction, informing our choice of the GP as surrogate model. As we will demonstrate in Section 5.3, active learning is a crucial component of the framework, ruling out a priori trained surrogate models.

Using this set of problems we investigate the following aspects of surrogate accelerated Bayesian inference: (i) How does the choice of data collection strategy impact the accuracy of the surrogate model? (ii) Do superior MCMC algorithms construct better surrogate models? (iii) After training the surrogate model, how does the sampling algorithm for the generation of posterior samples impact accuracy and performance?

2. Background

2.1. Computational mechanics

This section presents a brief introduction to computational mechanics. The deformation of the body within the domain \mathcal{D} with boundary \mathcal{B} is governed by the equilibrium equation, which can be written as

$$\nabla \cdot \sigma = 0 \quad \text{for } \mathbf{x} \in \mathcal{D}, \quad (1)$$

$$\mathbf{u}(\mathbf{x}, t) = \hat{\mathbf{u}}(\mathbf{x}, t) \quad \text{for } \mathbf{x} \in \mathcal{B}, \quad (2)$$

where σ is the stress tensor and inertia effects and external forces are assumed to be negligible. The boundary conditions in Eq. (2) enforce a displacement $\hat{\mathbf{u}}$ at the boundaries of the domain. The kinematics of the deformation are described by the displacement field $\mathbf{u}(\mathbf{x})$. Assuming small strains, the strain tensor ϵ is given by

$$\epsilon = \frac{1}{2} (\nabla \mathbf{u} + (\nabla \mathbf{u})^\top). \quad (3)$$

This kinematic relation must be linked to the equilibrium equation in Eq. (1). The constitutive model C relates the stresses to the strains and a local material parameter β :

$$\sigma = C(\epsilon, \beta). \quad (4)$$

Here, β denotes a (fixed) material parameter. It is distinct from internal, history-dependent state variables (e.g. plastic strain, damage) whose evolution is prescribed by the constitutive model. These states are therefore not inferred directly; their evolution laws are then fully dictated by β . Often, the local parameter value is expressed in terms of a set of global parameters $\theta \in \mathbb{R}^d$ and the location \mathbf{x} , e.g. $\beta = \beta(\mathbf{x}, \theta) = \theta_0 + \mathbf{x} \cdot \theta_1$.

Substituting Eq. (4) into Eq. (1) gives

$$\nabla \cdot (C(\epsilon, \beta)) = 0, \quad (5)$$

which together with Eq. (3) can be solved for the displacement field $\mathbf{u}(\mathbf{x})$ and yields the forward model $\mathcal{M}(\theta)$. While the extension to the multi-parameter case is straightforward, we will focus on the scalar case here for simplicity.

The choice of constitutive model — be it linear or non-linear, isotropic or anisotropic, and dependent or independent of time — hinges on the specific material and its regime of deformation. Irrespective of the clarity in the nature of the chosen constitutive model, a persistent challenge lies in the inherent uncertainty of material parameters (Rappel et al., 2019, 2020). These parameters, crucial for accurately predicting material behaviour, are often not directly measurable but must be inferred from experimental data. The task of inferring these parameters frames the inverse problem associated with the forward model \mathcal{M} .

2.2. Random fields and their discretisation

Acknowledging the lack of total control over manufacturing processes or the inherent randomness of the material itself, the assumption of homogeneous material properties is generally not justified. In many applications, material properties can vary significantly across different regions of a part and exhibit randomness under repeated production. We are, therefore, interested not just in inferring global material parameters but also how they vary throughout \mathcal{D} . However, there is typically some structure to the randomness that we need to consider when inferring the spatial distribution of material properties.

A useful way to capture this spatial variability is through random fields, which provide a framework for describing quantities that vary both in space and with some inherent randomness. More specifically, a random field is a collection of random variables associated with points in a physical space, $\mathbf{x} \in \mathcal{D} \subset \mathbb{R}^D$. For simplicity, we can think of a random field as a function $\hat{\beta}(\mathbf{x})$ that assigns random values to each point \mathbf{x} in the domain.

Many physical quantities, such as material properties, can be modelled as Gaussian random fields. This means that any set of values the field takes at different points, $\hat{\beta}(\mathbf{x}_1), \dots, \hat{\beta}(\mathbf{x}_k)$ for $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathcal{D}$, follows a multivariate normal distribution. This type of field is fully described by two key statistical properties: (i) the mean function $\mu_{\hat{\beta}}(\mathbf{x})$, which describes the average value of the field at each point, and (ii) the covariance function $C_{\hat{\beta}}(\mathbf{x}, \mathbf{x}')$, which tells us how values at different points are correlated based on their distance apart.

In practice, the field is expressed as a truncated series expansion in terms of a set of deterministic basis functions $\phi_i(\mathbf{x})$ with random coefficients θ_i . By cutting off the series expansion at a finite number of terms d , we obtain a finite-dimensional approximation of the random field:

$$\hat{\beta}(\mathbf{x}) \approx \beta(\mathbf{x}, \theta) = \mu_{\hat{\beta}}(\mathbf{x}) + \sum_{i=1}^d \phi_i(\mathbf{x}) \theta_i. \quad (6)$$

The expansion effectively amounts to a separation of the physical and the stochastic domain (Sudret and Kiureghian, 2000). Since all randomness is now contained in the coefficients $\theta \in \mathbb{R}^d$, inferring the random field reduces to inferring θ .

Expansions encountered in related works employ a basis of polynomials (Deveney et al., 2023; Marzouk and Najm, 2009), wavelets (Nouy and Maître, 2009), B-splines (Vigliotti et al., 2018), or Gaussian-process shape functions (Li and Kiureghian, 1993). The optimal expansion in terms of the total mean squared error of the truncation compared to the full sum is the Karhunen-Loève (KL) expansion. It is the eigenfunction expansion of the covariance function of the random field (Ghanem and Spanos, 1991). For a Gaussian random field, the KL expansion coefficients θ are i.i.d. standard normal random variables. For the other expansions, the coefficients are normally distributed but correlated.

We opt for a radial basis function (RBF) expansion. The basis is given by

$$\phi_i(\mathbf{x}) = \exp\left(-\frac{1}{2l^2} \|\mathbf{x} - \mathbf{c}_i\|^2\right), \quad (7)$$

where \mathbf{c}_i are the centres of the RBFs and l is the length-scale. Each basis function ϕ_i reaches its maximum at \mathbf{c}_i and decays with increasing distance from the centre. This choice is not optimal in terms of the total mean squared error for a given number of basis functions. However, it has a clear advantage for defining a sequence of comparable inference problems: the RBF expansion leads to a parameterisation of the field where all coefficients have a similar influence on the global system response. On the other hand, the ordering of the eigenfunctions of the KL expansion according to their eigenvalues ensures that the model is most sensitive to changes in the first few coefficients. Furthermore, we circumvent the need to solve the associated eigenvalue problem.

Given that the random field is characterised by a squared exponential covariance function, employing a linear model with RBFs and appropriately setting the RBF length-scale produces a similar covariance structure. It converges to the desired covariance in the limit of infinitely many basis functions (Rasmussen and Williams, 2006, p. 84). Note that an RBF model with length-scale l produces a covariance function with effective length-scale $\ell = \sqrt{2}l$ (Mackay, 1998). For simplicity, we choose the prior of the weights to be i.i.d Gaussian, i.e. $\theta \sim \mathcal{N}(\mathbf{0}, \sigma_\theta^2 \mathbf{I})$. It is convenient to store the basis functions in a feature vector $\phi(\mathbf{x}) = [\phi_1(\mathbf{x}), \dots, \phi_d(\mathbf{x})]^\top$. The prior mean of the RBF expansion then reads:

$$\mathbb{E}[\beta(\mathbf{x}, \theta)] = \mathbb{E}[\mu_{\hat{\beta}}(\mathbf{x}) + \phi(\mathbf{x})^\top \theta] = \mu_{\hat{\beta}}(\mathbf{x}) + \phi(\mathbf{x})^\top \mathbb{E}[\theta] = \mu_{\hat{\beta}}(\mathbf{x}), \quad (8)$$

showing the desired mean can be achieved by setting $\mu_{\hat{\beta}} = \mu_{\hat{\beta}}$. The covariance of the RBF expansion is given by:

$$\text{cov}(\beta(\mathbf{x}, \theta), \beta(\mathbf{x}', \theta)) = \mathbb{E}[(\beta(\mathbf{x}, \theta) - \mu_{\hat{\beta}}(\mathbf{x}))(\beta(\mathbf{x}', \theta) - \mu_{\hat{\beta}}(\mathbf{x}'))] \quad (9)$$

$$= \mathbb{E}[\phi(\mathbf{x})^\top \theta \theta^\top \phi(\mathbf{x}')] = \phi(\mathbf{x})^\top \phi(\mathbf{x}') \sigma_\theta^2. \quad (10)$$

The variance consequently is

$$\sigma_{\hat{\beta}}^2(\mathbf{x}) = \phi(\mathbf{x})^\top \phi(\mathbf{x}) \sigma_\theta^2. \quad (11)$$

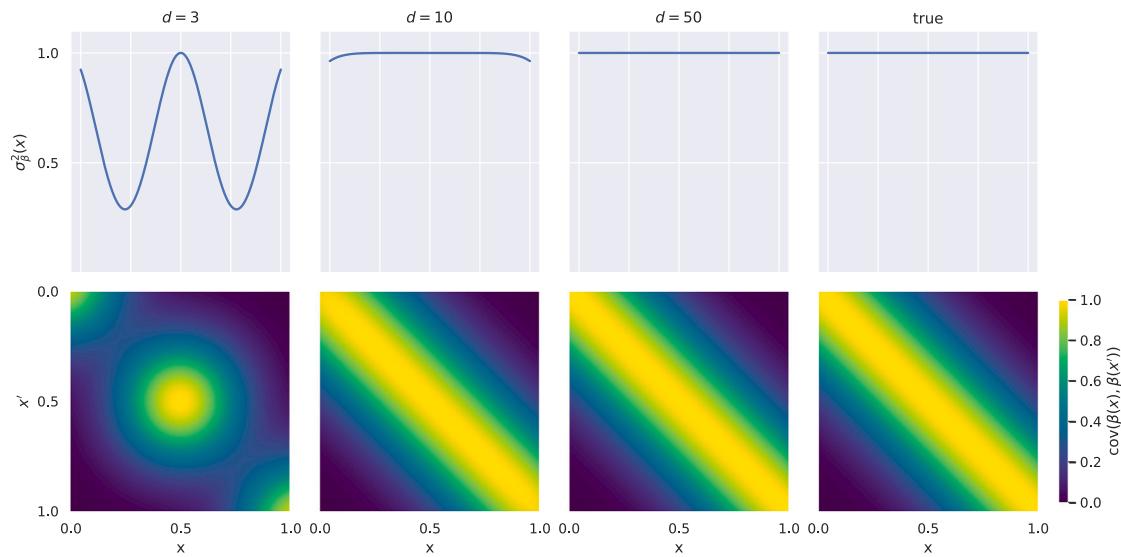


Fig. 1. Discretisation of the random field with increasing amount of radial basis functions. The variance of the fields is shown in the top row. The banded covariance structure is displayed in the bottom row.

The prior variance of the weights σ_θ^2 can be set to match the variance of the random field of interest. It is then given by

$$\sigma_\theta^2 = \frac{\sigma_\beta^2}{\sup \{ \phi(x)^\top \phi(x), x \in \mathcal{D} \}}. \quad (12)$$

where the denominator scales the maximum variance of the RBF expansion to the desired variance of the random field.

The random field parameters, i.e. its mean μ_β , variance σ_β^2 , and effective length-scale ℓ , can be inferred along the coefficients θ in a hierarchical Bayesian model. We assume to have prior knowledge of the problem that we can use to set these hyperparameters, as hierarchical modelling is not the focus of this study.

Fig. 1 shows the discretisation of a random field on a one-dimensional domain \mathcal{D} with varying amounts of basis functions. As the number of basis functions increases, the variance of the field approaches a constant value throughout the domain. The banded covariance structure on the far right of the figure is well approximated by the RBF model. We can see that the expansion produces a stationary random field that converges in variance and length-scale to the infinite-dimensional random field.

2.3. Bayesian inference for random fields

Our PDE model \mathcal{M} maps material parameters $\theta \in \mathbb{R}^d$ to an output $y \in \mathbb{R}^m$, i.e. $y = \mathcal{M}(\beta(\cdot, \theta)) = \mathcal{M}(\theta)$. The inverse problem is concerned with inferring these model parameters from observation data \tilde{y} . In mechanics, the observation data \tilde{y} typically comprises noisy measurements of the displacement field u and the reaction forces f , i.e. $\tilde{y} = [u^\top, f^\top]^\top$. For a Bayesian approach to inverse problems, the parameters and the data are represented as random vectors and equipped with a probability measure, see [Stuart \(2010\)](#) for a detailed description. The prior distribution for the parameters θ prior to seeing any data. The likelihood function $p(\tilde{y}|\theta)$ expresses how well the model prediction matches the observed data for a given realisation of the parameter vector θ . In our case, the likelihood is a function of the model \mathcal{M} and takes the form

$$p(\tilde{y}|\theta) = p(\tilde{y}|\mathcal{M}(\theta)), \quad (13)$$

and is induced by the data-generation mechanism: measurements equal the model response plus noise. We assume additive noise that is independent of θ , capturing the uncertainty in \tilde{y} :

$$\tilde{y} = \mathcal{M}(\theta) + \epsilon \quad (14)$$

where $\epsilon \in \mathbb{R}^m$ is the noise random vector. In this work, the noise vector is assumed to follow a zero-mean Gaussian distribution with covariance matrix $\Sigma_{\epsilon\epsilon} \in \mathbb{R}^{m \times m}$. More specifically, each data point comprises displacement and force observations f and u , respectively, which come from measuring devices with distinct noise levels for ϵ_u and ϵ_f . Extending Eq. (14) to the multi-output case, we have

$$\tilde{u} = u + \epsilon_u, \quad \tilde{f} = f + \epsilon_f. \quad (15)$$

Assuming no correlation between the two types of observations, their joint likelihood takes the form

$$p\left(\begin{bmatrix} u \\ f \end{bmatrix} \mid \theta\right) = \mathcal{N}\left(\begin{bmatrix} u \\ f \end{bmatrix} \mid \mathcal{M}(\theta), \begin{bmatrix} \Sigma_{uu} & \mathbf{0} \\ \mathbf{0} & \Sigma_{ff} \end{bmatrix}\right), \quad (16)$$

where the covariance structures Σ_{uu} and Σ_{ff} of the displacement and force observations, respectively, are known a priori.

It should be noted that the i.i.d. noise model choice is based on the assumption that we exactly know the physics of the problem, i.e. there is no model misspecification, and that our measurement devices show no spatial or temporal correlation. This simplification is justified for the scope of this study, as it investigates aspects of the Bayesian inference workflow other than the definition of the likelihood function. In practice, however, a more complex observation model that reflects all sources of uncertainty, as outlined in [Simoen et al. \(2013\)](#), should be employed. Otherwise, the inference might yield overconfident posterior beliefs about the parameters governing the random field, which can be misleading in a subsequent decision-making process.

Finally, we can apply Bayes' theorem to obtain the density of the parameters θ conditioned on the observations \tilde{y} , which reads

$$\underbrace{p(\theta|\tilde{y})}_{\text{posterior}} = \frac{\underbrace{p(\tilde{y}|\theta)}_{\text{likelihood}} \underbrace{p(\theta)}_{\text{prior}}}{\underbrace{p(\tilde{y})}_{\text{model evidence}}}, \quad (17)$$

where the denominator on the right hand side is the so-called evidence. This normalising constant can be obtained by marginalising out the parameters:

$$p(\tilde{y}) = \int_{\mathbb{R}^d} p(\tilde{y}|\theta)p(\theta)d\theta. \quad (18)$$

With a few exceptions, the integral in Eq. (18) is intractable. While the evidence can be approximated via numerical methods such as Monte Carlo simulation, it is notoriously hard to compute accurately.

More efficiently, we can draw samples directly from the unnormalised posterior distribution via MCMC methods. For the latter approach, we only need to know the posterior distribution up to a proportionality constant, hence no need to obtain the evidence. We can draw samples from the posterior despite only having access to the joint distribution of data and parameters:

$$p(\theta|\bar{y}) \propto p(\bar{y}, \theta) = p(\bar{y}|\theta)p(\theta). \quad (19)$$

2.4. Markov chain Monte Carlo methods

MCMC methods are a class of algorithms used to sample from a probability distribution from which direct sampling is difficult, hence they are suitable for Bayesian inference in computational mechanics. A Markov chain that has the desired distribution as its equilibrium distribution is constructed and its states are recorded. The more steps the chain takes, the more the distribution of the states of the chain converges to the target distribution.

2.4.1. Random walk Metropolis algorithm

The RWM algorithm is based on a random walk in the state space, determined by a proposal distribution to generate new states. In the subsequent Metropolis correction, the proposed state is either accepted so that the chain moves, or rejected so that the chain stays in place. This correction is essential for the convergence of the chain to the target distribution and is a core building block for many other MCMC algorithms.

Suppose we have a function $f(\theta)$ that is proportional to our target distribution $p(\theta|\bar{y})$. In our case, this function is the joint distribution of parameters and data $p(\bar{y}, \theta)$. Once the data is observed, the joint distribution essentially becomes a function of the parameters θ only. Given the current state θ^n , a move to θ^* is proposed according to the proposal distribution $q(\theta^*|\theta^n)$. Next, the Metropolis–Hastings ratio is calculated as

$$r(\theta^n, \theta^*) = \frac{f(\theta^*)q(\theta^*|\theta^n)}{f(\theta^n)q(\theta^n|\theta^*)}. \quad (20)$$

The move from θ^n to θ^* is then accepted with probability

$$a(\theta^n, \theta^*) = \min(1, r(\theta^n, \theta^*)). \quad (21)$$

If the proposal distribution $q(\theta^*, \theta^n)$ is symmetric — the normal distribution is a common choice — the ratio simplifies to the Metropolis ratio

$$r(\theta^n, \theta^*) = \frac{f(\theta^*)}{f(\theta^n)}, \quad (22)$$

also reflected in the algorithms name.

The Gaussian proposal mechanism can then be written as

$$\theta^* = \theta^n + s\sqrt{\mathbf{M}}\mathbf{z}^n, \quad (23)$$

where \mathbf{M} is a positive semi-definite preconditioner matrix which defines the covariance of the proposal distribution, \mathbf{z}^n is a vector of i.i.d. standard normal random variables, and s is the proposal variance controlling the step size. The corresponding proposal distribution then is

$$q(\theta^*|\theta^n) = \mathcal{N}(\theta^*|\theta^n, s^2\mathbf{M}). \quad (24)$$

The performance of the algorithm can be significantly improved by setting \mathbf{M} to the sample covariance. The sample covariance can either be estimated from a preliminary run of the algorithm or with a Laplace approximation of the posterior. The proposal covariance can also be adapted continuously during runtime, as long as the level of adaptation is vanishing with the number of iterations (Roberts and Rosenthal, 2007). Because diminishing adaptiveness is difficult to prove in practice, we set the proposal covariance to the sample covariance only once after the first 5000 MCMC-steps. Until then, we use the prior

covariance as preconditioner. This version of the RWM algorithm serves as the baseline MCMC algorithm in our study.

High acceptance rates can be achieved by reducing the variance of the proposal distribution, but this can lead to slow exploration of the state space. Conversely, a high variance can lead to fast exploration, but with a low acceptance rate. It is, therefore, important to choose the step size s in a way that balances these two aspects. However, Gelman et al. (1997) have shown that the proposal variance s must be scaled with the state space dimensionality d according to $\mathcal{O}(d^{-1})$ to achieve the optimal acceptance rate of 0.234. Hence, if d is large, the transitions are small and the algorithm shows poor mixing and slow convergence. Nonetheless, this poor scaling with d is one of the main motivations to opt for gradient-based MCMC algorithms, such as the Metropolis-adjusted Langevin algorithm (MALA).

We constantly monitor the acceptance rate and adjust the step size s to achieve the algorithm's optimal acceptance rate. To this end, the acceptance rate of the previous interval is determined every 500 MCMC steps. The step size s is then either increased or decreased by 10%, if the acceptance rate is below 0.2 or above 0.3, respectively. If the rate is within the desired range, the step size is kept constant.

2.4.2. Metropolis-adjusted Langevin algorithm

The Langevin algorithm (Roberts and Stramer, 2002) derives its proposal mechanism from a discretised Langevin diffusion. It can be seen as a random walk with a drift term that is proportional to the gradient of the log-density. While the continuous Langevin diffusion has the targeted log-density as its stationary distribution, its discretisation introduces a bias. The Metropolis-adjusted Langevin algorithm (MALA) corrects for this bias by adding a Metropolis–Hastings correction to the mechanism.

The proposal mechanism of the MALA is given by

$$\theta^* = \theta^n + \frac{s^2}{2} \mathbf{M} \nabla \ln f(\theta^n) + s\sqrt{\mathbf{M}}\mathbf{z}^n. \quad (25)$$

which gives rise to the Gaussian proposal distribution

$$q(\theta^*|\theta^n) = \mathcal{N}\left(\theta^*|\theta^n + \frac{s^2}{2} \mathbf{M} \nabla \ln f(\theta^n), s^2 \mathbf{M}\right). \quad (26)$$

The MALA proposal is similar to the RWM proposal Eq. (24), but has the additional drift term which is proportional to the gradient of the unnormalised log-density $\ln f(\theta^n)$. Note that proposal distribution is non-symmetric in θ^n and θ^* . Hence, the Metropolis–Hastings ratio does not simplify to the ratio of the densities and has to be computed according to the general formula in Eq. (20).

As for the RWM algorithm, the preconditioner matrix \mathbf{M} is set to the sample covariance after the first 5000 MCMC-steps. Similarly, the step size s is adaptively tuned during runtime to achieve the optimal acceptance rate of 0.574 (Roberts and Rosenthal, 1998). Only, the target interval for the acceptance rate is set to [0.55, 0.6] for the MALA. The MALA is known to be more efficient than the RWM algorithm for high-dimensional state spaces, as it can exploit the gradient information to make larger steps in the direction of the mode of the target distribution.

2.5. Gaussian process surrogate model

We employ a GP surrogate model to approximate the likelihood function $\mathcal{L}(\theta)$ in order to reduce the computational cost associated with likelihood evaluations in the many MCMC iterations. From a mathematical angle, a GP is a collection of random variables, any finite number of which have a joint Gaussian distribution:

$$\mathcal{L} \sim \mathcal{N}(\mathcal{L}|\mathbf{0}, \mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta})), \quad (27)$$

where a zero mean is assumed before seeing any data. The covariance matrix $\mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta})$ is the collection of the kernel function on all input pairs, i.e. $K_{ij} = k(\theta_i, \theta_j)$ The matrix specifies the similarities in the output \mathcal{L}

based on their input $\boldsymbol{\theta}$. The kernel in this case is chosen as the squared exponential kernel (Rasmussen and Williams, 2006), given by

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_f^2 \exp\left(-\frac{1}{2\ell^2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2\right). \quad (28)$$

The GP variance σ_f^2 controls the magnitude of the fluctuations around the mean function, and its length scale ℓ controls the smoothness of the function. We assume the relationship of input and output of our training data to be of the form $\boldsymbol{\eta}_i = \mathcal{L}(\boldsymbol{\theta}_i) + \boldsymbol{\epsilon}_i$, where $\mathcal{L}(\boldsymbol{\theta})$ is the likelihood function and $\boldsymbol{\epsilon}_i$ is i.i.d. noise with variance σ_n^2 . Although the training data originates from noise-free simulations, we admit a noise term $\sigma_n^2 \mathbf{I}$ for numerical stability and model flexibility. We denote $\mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta})$ the covariance matrix between the training and test points, and $\mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*)$ the covariance matrix between the test points. We can then write the joint distribution over the training output $\boldsymbol{\eta}$ and new predictions \mathcal{L}^* at locations $\boldsymbol{\theta}^*$ as

$$\begin{bmatrix} \boldsymbol{\eta} \\ \mathcal{L}^* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta}) + \sigma_n^2 \mathbf{I} & \mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta}^*) \\ \mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}) & \mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) \end{bmatrix}\right). \quad (29)$$

Applying Bayes' formula, the predictive distribution of the test output $\mathcal{L}^* = \mathcal{L}(\boldsymbol{\theta}^*)$ given the training data $\mathcal{D}_{\text{GP}} = (\boldsymbol{\theta}, \boldsymbol{\eta})$ is then given by

$$\begin{aligned} \mathcal{L}^* | \boldsymbol{\eta}, \boldsymbol{\theta}, \boldsymbol{\theta}^* &\sim \mathcal{N}(\mathcal{L}^* | \boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*), \\ \boldsymbol{\mu}^* &= \mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}) (\mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta}) + \sigma_n^2 \mathbf{I})^{-1} \boldsymbol{\eta}, \\ \boldsymbol{\Sigma}^* &= \mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) - \mathbf{K}(\boldsymbol{\theta}^*, \boldsymbol{\theta}) (\mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta}) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{K}(\boldsymbol{\theta}, \boldsymbol{\theta}^*). \end{aligned} \quad (30)$$

The predictive mean $\boldsymbol{\mu}^*$ is our best guess about the function values at the test points, and the predictive variance $\boldsymbol{\Sigma}^*$ quantifies our uncertainty about these predictions.

We determine the GP model hyperparameters — the process variance σ_f^2 , the noise variance σ_n^2 , and the length scale ℓ — via empirical Bayes for computational efficiency. For this, we maximise the marginal likelihood of the training outputs $\boldsymbol{\eta}$ given the training input $\boldsymbol{\theta}$. The marginal likelihood $p(\boldsymbol{\eta} | \sigma_f^2, \sigma_n^2, \ell)$ and its logarithm L are given by

$$p(\boldsymbol{\eta} | \sigma_f^2, \sigma_n^2, \ell) = \int p(\boldsymbol{\eta} | \mathcal{L}, \sigma_n^2) p(\mathcal{L} | \sigma_f^2, \ell) d\mathcal{L}, \quad (31)$$

$$L = \ln p(\boldsymbol{\eta} | \sigma_f^2, \sigma_n^2, \ell). \quad (32)$$

Instead of maximising the marginal likelihood directly, we maximise the log marginal likelihood Eq. (32) to avoid numerical issues. The associated optimisation problem is solved with the BFGS algorithm (Fletcher, 2000).

3. MCMC-guided active learning

In this section, we introduce an algorithm that employs MCMC-guided sampling as a straightforward active learning strategy for constructing surrogate models. This approach leverages the predictive uncertainty inherent in the GP surrogate model, coupled with the trajectory of the MCMC algorithm, drawing inspiration from the active learning framework presented in Rocha et al. (2021).

Initially, we generate a set of N_0 samples from the posterior distribution using the forward model. This set forms the initial training dataset for our GP surrogate model. Upon this dataset, we fit the GP model by estimating its hyperparameters and record the initial log marginal likelihood, denoted as L_{old} . After the initialisation phase is completed, the MCMC algorithm generates new proposals based on the GP model. Denoting $\mathbf{k}(\boldsymbol{\theta}^*)$ the covariance vector between the training data and the new point $\boldsymbol{\theta}^*$, and \mathbf{K} the covariance matrix of the training data, the predictive variance of the GP model at $\boldsymbol{\theta}^*$ is given by

$$\mathbb{V}[\mathcal{L}(\boldsymbol{\theta}^*)] = \mathbf{k}(\boldsymbol{\theta}^*, \boldsymbol{\theta}^*) - \mathbf{k}(\boldsymbol{\theta}^*)^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}(\boldsymbol{\theta}^*). \quad (33)$$

If this variance remains below a specified threshold γ_v , we accept the GP's prediction at $\boldsymbol{\theta}^*$ and proceed without further action. Conversely, if the predictive variance exceeds the variance threshold γ_v , we resort to evaluating the forward model $\mathbf{y}^* = \mathcal{M}(\boldsymbol{\theta}^*)$ and compute the likelihood

$\mathcal{L}(\boldsymbol{\theta}^*)$ to obtain a new observation. We will refer to γ_v as the reject threshold. This observation is then used to update our training dataset.

Following the addition of new data, we update the GP's posterior covariance matrix and recalculate the log marginal likelihood, denoted as L_{new} . If the ratio $|L_{\text{new}}/L_{\text{old}}|$ exceeds a predetermined threshold γ_L , this indicates a significant change in the model's understanding, prompting a re-estimation of the GP hyperparameters based on the updated dataset. We will thus refer to γ_L as retrain threshold. The process iterates until the end of the burn-in phase N_b is reached, at which point the GP model's hyperparameters are fixed. These parameters have a significant impact on the model's predictive accuracy and their choice is therefore investigated in Section 5.2.

Replacing the forward model with the GP surrogate model leads to a modified proposal distribution Eq. (26) for the MALA:

$$\tilde{q}(\boldsymbol{\theta}^* | \boldsymbol{\theta}^n) = \mathcal{N}\left(\boldsymbol{\theta}^* | \boldsymbol{\theta}^n + \frac{s^2}{2} \mathbf{M}^{-1} \nabla_{\boldsymbol{\theta}} (\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^n)] + \ln p(\boldsymbol{\theta}^n)), s^2 \mathbf{M}\right). \quad (34)$$

Here, the gradient of the log-likelihood is computed using the mean of the GP posterior $\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^n)]$. We also need to add the contribution of the log-prior $\ln p(\boldsymbol{\theta}^n)$. The modified MALA acceptance probability is then given by:

$$\tilde{a}(\boldsymbol{\theta}^n, \boldsymbol{\theta}^*) = \min\left(1, \frac{\exp(\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^*)]) p(\boldsymbol{\theta}^*) \tilde{q}(\boldsymbol{\theta}^* | \boldsymbol{\theta}^n)}{\exp(\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^n)]) p(\boldsymbol{\theta}^n) \tilde{q}(\boldsymbol{\theta}^n | \boldsymbol{\theta}^*)}\right). \quad (35)$$

The proposal for the RWM algorithm Eq. (24) remains the same, as it does not depend on the forward model. However, the acceptance probability Eq. (21) is modified to account for the GP surrogate model:

$$\tilde{a}(\boldsymbol{\theta}^n, \boldsymbol{\theta}^*) = \min\left(1, \frac{\exp(\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^*)]) p(\boldsymbol{\theta}^*)}{\exp(\mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^n)]) p(\boldsymbol{\theta}^n)}\right). \quad (36)$$

Note that we need to exponentiate the GP prediction in Eqs. (34) and (36), as we trained the GP on the log-likelihood. The MCMC algorithm incurs a bias from the surrogate model in the acceptance probabilities Eq. (35) and (36). The quantification of this bias and its impact on the posterior estimation is at the heart of this work.

As the GP model refines its accuracy with the incorporation of new data, we anticipate a reduction in predictive variance, signalling an enhanced model fidelity. Consequently, the need to resort to the computationally expensive forward model diminishes. This strategy ensures that new data points are strategically added to the model's training set, particularly focusing on regions yet to be explored by the MCMC algorithm, without changing the response surface in regions that have been visited in previous steps. The active learning strategy is described in detail in Algorithm 1.

4. Case study

In this section, we introduce a scalable case study to elucidate the effects of the methodological decision on surrogate modelling and sampling. First and foremost, the problem must be computationally inexpensive. Since we cannot obtain an analytical posterior, we need the capability to brute-force approximate this distribution via long runs of MCMC to obtain a reference solution. Additionally, the problem should closely resemble a real-world scenario. To this end, we introduce a one-dimensional bar with length $L = 1 \text{ m}$ and constant cross-sectional area $A = 1 \text{ m}^2$, depicted in figure Fig. 2.

4.1. Material model

To make things more challenging, we introduce a simple non-linear constitutive model and a spatially varying initial stiffness. It is adopted as a fast and pragmatic choice, but nevertheless without loss of generality. The constitutive model is given by

$$\sigma(\varepsilon) = r(E - H) \left(1 - \exp\left(-\frac{\varepsilon}{r}\right)\right) + H\varepsilon. \quad (37)$$

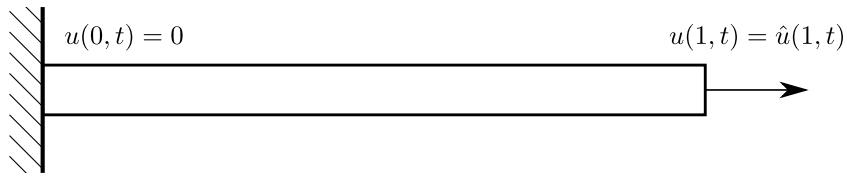


Fig. 2. Schematic representation of the one-dimensional bar problem. Dirichlet boundary conditions are enforced at the left and right boundaries.

Algorithm 1 MCMC-guided active learning for GP surrogate model construction. Depending on the choice of the algorithm, the proposal density $\tilde{q}(\theta^*|\theta^n)$ in Line 5 is either given by Eq. (24) for the RWM algorithm or Eq. (34) for the MALA. Similarly, the acceptance probability α in Line 17 is given by Eq. (36) (RWM) or Eq. (35) (MALA).

```

1: Initialise  $N_0$  samples from the posterior using the forward model to
   obtain initial dataset  $D_{\text{GP}} = \{(\theta_0, \mathcal{L}(\theta_0)), \dots, (\theta_{N_0}, \mathcal{L}(\theta_{N_0}))\}$ 
2: Estimate hyperparameters on  $D_{\text{GP}}$  and record  $L_{\text{old}}$ 
3:  $n \leftarrow 0$ 
4: while  $n < N_{\text{total}}$  do
5:   Generate new point  $\theta^*$  from the proposal distribution  $\tilde{q}(\theta^*|\theta^n)$ 
6:   if  $\nabla[\mathcal{L}(\theta^*)] < \gamma_v$  then
7:     Accept GP prediction for  $\mathcal{L}(\theta^*)$ 
8:   else
9:     Evaluate forward model  $y^* = \mathcal{M}(\theta^*)$ , compute likelihood
    $\mathcal{L}(\theta^*)$ 
10:    Update training dataset  $D_{\text{GP}} \leftarrow D_{\text{GP}} \cup (\theta^*, \mathcal{L}(\theta^*))$ 
11:    Update GP's posterior covariance matrix and compute  $L_{\text{new}}$ 
12:    if  $|L_{\text{new}}/L_{\text{old}}| > \gamma_L$  AND  $n < N_b$  then
13:      Re-estimate GP hyperparameters and update  $L_{\text{new}}$ 
14:       $L_{\text{old}} \leftarrow L_{\text{new}}$ 
15:    end if
16:  end if
17:  if  $\tilde{\alpha}(\theta^n, \theta^*) > u \sim \mathcal{U}(0, 1)$  then
18:     $\theta^{n+1} \leftarrow \theta^*$ 
19:  else
20:     $\theta^{n+1} \leftarrow \theta^n$ 
21:  end if
22:   $n \leftarrow n + 1$ 
23: end while

```

This model, characterised by its initial stiffness E , its terminal stiffness H , and its rate parameter r , mimics nonlinear material behaviour during plastic deformation. It does so without necessitating the computational burden associated with tracking the evolution of plastic strain. Note that this simplification is only reasonable for the monotonic loading considered in this study. We thereby embed additional non-linearity into the system without the exhaustive computational requirements typical of a complete plasticity model.

Fig. 3(a) illustrates the stress-strain relationship for different values of the material parameters. The material model's non-linearity is evident, with the stress-strain curve exhibiting a distinct hardening behaviour. After initial numerical testing, no single parameter stood out as comparatively more challenging to infer, except for a hardening modulus H close to zero. We, therefore, focus exclusively on the initial stiffness E in the subsequent analysis to ensure a consistent comparison across the different methods at fixed H and r . The spatially varying initial stiffness $\hat{\beta} = \hat{E} = \hat{E}(x) \approx E(x, \theta)$ is hence the unknown field to be inferred.

4.2. Sequence of inverse problems

We define a sequence of inverse problems by considering different stochastic discretisations of the spatially varying initial stiffness. Taking the spatial correlation of the RBF discretisation into account, further

additions of basis functions are expected to yield diminishing gains of information for a given random field. As the RBFs with centres close to one another are expected to correlate, their inference will be less challenging. We adapt the length scale of the random field, and thereby the complexity of the inference problem, according to the number of terms in the random field expansion, as demonstrated in Fig. 4. In other words, we must infer a random function with an expected number of zero crossings proportional to the number of free parameters. We also increase the amount of displacement data to ensure that the observations-to-parameters ratio remains constant. This ensures any difference in inference performance as the dimensionality of the latent space increases can be ascribed to the dimensionality and not to the complexity of the inferred behaviour.

To enforce positivity of the stiffness field, we parameterise it via a log-field. Let $\hat{E}(x)$ denote the Gaussian log-stiffness field and $\bar{E}(x, \theta)$ its RBF discretisation. We then define the corresponding stiffness fields by exponentiation, $\hat{E}(x) = \exp(\hat{E}(x))$ and $E(x, \theta) = \exp(\bar{E}(x, \theta))$ (in MPa). Consequently, if $\hat{E}(x)$ is Gaussian, $\hat{E}(x)$ is log-normal pointwise. The explicit relations are:

$$\hat{E}(x) = \exp(\hat{E}(x)), \quad (38)$$

$$E(x, \theta) = \exp(\bar{E}(x, \theta)) = \exp\left(\mu_{\bar{E}}(x) + \sum_{i=1}^d \phi_i(x) \theta_i\right), \quad (39)$$

where $\mu_{\bar{E}}(x)$ is the mean of the underlying normal distribution. We set the random field parameters to $\mu_{\bar{E}}(x) = 8$, $\sigma_{\theta} = 0.1$ to mimic an elastoplastic polymer used as matrix material in fibre reinforced composites (Melro et al., 2013), and fix $l = 1.5/d$. The corresponding marginal distribution is depicted in Fig. 3(b) and the random field realisations for different numbers of RBFs and respective length scales are shown in Fig. 4. The goal is then to find the posterior distribution of the coefficients θ of the underlying Gaussian random field, which can then be transformed to the posterior of the parameter field with help of Eq. (39). We fix the remaining parameters to $H = 100$ MPa and $r = 0.01$.

In our fictitious experiment, we enforce Dirichlet boundary conditions on the left and right boundaries as follows:

$$u(0, t) = 0, \quad (40)$$

$$u(1, t) = t \quad \text{for } t \in [0, T]. \quad (41)$$

The displacement of the right boundary increases linearly with pseudo time t and the final time is set to $T = 0.1$ s. We sample the ground truth from the prior distribution (see Eq. (39) and Fig. 4) and solve the forward problem. To generate the observations, we record the displacements at a number of locations and the reaction force at the far right boundary at $N_j = 5$ distinct times $t_j \in \{0.02, 0.04, 0.06, 0.08, 0.1\}$ s. The number of recorded displacements per time t_j is set to $N_i = \lceil 3d/4 \rceil$, where $\lceil \cdot \rceil$ denotes the closest integer. This means there are fewer measurement locations x_i than parameters to infer. These locations are spread uniformly within the domain \mathcal{D} .

We then synthesise measurements from the recorded displacements and forces according to our observation model Eq. . We assume uncorrelated Gaussian noise, i.e. $\Sigma_{\mathbf{u}} = \sigma_u^2 \mathbf{I}$ and $\Sigma_{\mathbf{f}} = \sigma_f^2 \mathbf{I}$, with respective standard deviations $\sigma_u = 0.001$ m and $\sigma_f = 1$ MN. This virtual experiment is conducted only once per given ground truth. The displacements and forces are stacked, i.e. $\mathbf{u} = [u(x_0, t_0), u(x_1, t_0), \dots, u(x_{N_i}, t_{N_i})]$ and $\mathbf{f} =$

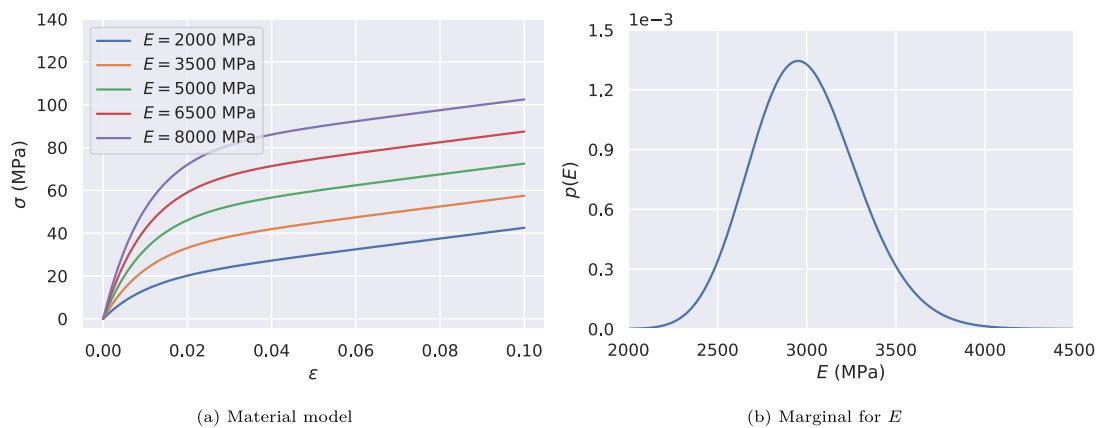


Fig. 3. (a) Non-linear material model with parameters $H = 250$ MPa and $r = 0.01$. (b) Marginal distribution of the log-normal field $E(x)$.

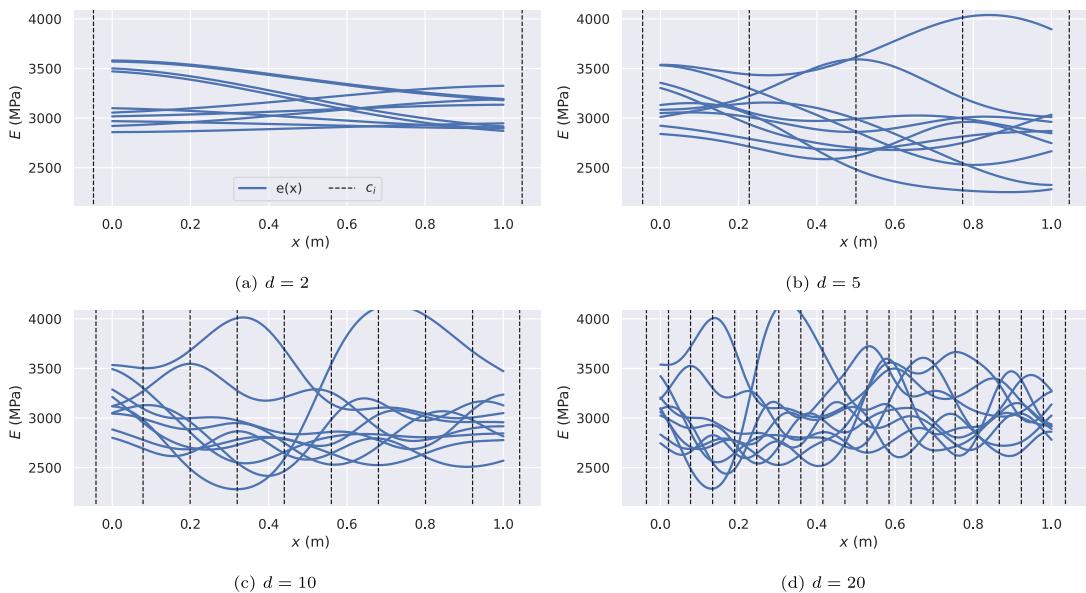


Fig. 4. Random field realisations of the spatially varying initial stiffness for different numbers of RBFs and respective length scales. The blue lines depict different generations of the ground truth, which is sampled from the prior distribution. The dashed black lines represent the centres of the RBFs.

$[f_0, \dots, f_{N_j}]$. The full observation vector $\tilde{y} = [\mathbf{u}^\top, \mathbf{f}^\top]^\top$ then comprises a total of $m = N_j \cdot (N_i + 1)$ scalar observations. The goal is now to infer the parameters of a matching discretisation of the random field from these noisy observations.

Remark. The actual stiffness field may not lie in the RBF span used for inference in practice, a situation which we do not treat systematically here. When the mismatch is small, Bayesian inference remains robust: the posterior balances data fit with prior regularisation and concentrates near the best in-span approximation. Practical remedies include a full Bayesian approach with a shrinkage prior on θ (Green, 1995), Bayesian model selection to choose the basis dimension (Uribe et al., 2020), or fitting a model discrepancy term (Kennedy and O'Hagan, 2001). However, these strategies come at a high computational cost and remain beyond the scope of this work.

4.3. Wasserstein distance

So far, we have introduced a scalable case study, two MCMC algorithms (RWM and MALA) to solve the inverse problems, and a GP surrogate model that can facilitate this process. To understand the impact of the methodological decisions on the inference process, we must compare the resulting approximate posterior distributions.

However, neither our reference posterior nor the approximation from the surrogate-model-assisted inference are available in analytical form, but instead are represented with samples from MCMC simulations. We must therefore find a suitable metric to compare these empirical distributions. The Wasserstein distance between two distributions is the minimum effort required to morph one into the other. This metric is particularly advantageous for empirical distributions, such as those represented by MCMC samples, because it can be directly applied to point clouds without requiring the distributions to be expressed in a functional form.

Other metrics, such as the Kullback–Leibler divergence or the total variation distance inherently require at least one of the distributions to be functionally described. To use these metrics in the given settings, one must approximate the empirical distribution using kernel density estimation or another density estimation technique, which introduces additional complexity and uncertainty into the analysis. In contrast, the Wasserstein distance directly considers the actual distances between points avoiding the pitfalls associated with density estimation. For our case study, we specifically employ the Wasserstein 2-distance, which is the square root of the optimal transport cost, and denote it as W_2 in the following.

The Wasserstein 2-distance $W_2(P, Q)$ between two continuous distributions P and Q on a metric space \mathcal{X} is defined as:

$$W_2(P, Q) = \left(\inf_{\gamma \in \Gamma(P, Q)} \int_{\mathcal{X} \times \mathcal{X}} \|x - y\|^2 d\gamma(x, y) \right)^{\frac{1}{2}}, \quad (42)$$

where $\Gamma(P, Q)$ is the set of all joint distributions on $\mathcal{X} \times \mathcal{X}$ with marginals P and Q . In case of empirical distributions represented by an equal amount of samples N_{total} , the Wasserstein 2-distance reduces to:

$$W_2(P, Q) = \left(\frac{1}{N_{\text{total}}} \sum_{i=1}^{N_{\text{total}}} \|x_i - y_{\pi(i)}\|^2 \right)^{\frac{1}{2}}, \quad (43)$$

where x_i and y_i are the samples from the distributions P and Q , respectively, and π is the permutation that minimises the square distance.

5. Results

In this section, we present the results for the case study. We first demonstrate the inference on the bar problem in general. Next, we validate the active learning approach on the five-dimensional problem and compare it to the a priori training strategies. Finally, we look at the impact of the MCMC algorithm on the surrogate model construction and on the efficiency and accuracy of the inference.

5.1. Reference solution

We first demonstrate the inference process on the one-dimensional bar problem with varying dimensionality of the unknown field. Fig. 5 shows the ground truth — a sample from the prior — alongside the prior and posterior distributions of the initial stiffness. All results are obtained using the RWM algorithm exclusively relying on the FEM model throughout the MCMC run. In these and all other runs, a total of 200000 MCMC samples were generated for each run, with the first 5000 samples discarded as burn-in. Further, each chain was thinned by a factor of 40 to ensure uncorrelated posterior samples. This yields a total of 4500 samples per posterior distribution, keeping the cost of computing the Wasserstein distance manageable. The posterior distributions of the initial stiffness are in good agreement with the ground truth in all scenarios presented in Fig. 5, underscoring the efficacy of the Bayesian approach. Further, the true fields are well contained within the 95% credible interval of the posterior. This consistency indicates a robust and meaningful quantification of the uncertainty, affirming the reliability of our inference process. With this non-linear, yet scalable sequence of inverse problems at hand, we now turn our attention to the many modelling choices in surrogate model assisted inference.

5.2. Active learning validation

Before diving into the comparison of the various training strategies, we must first validate the active learning approach.

5.2.1. Influence of active learning parameters

We use the 5-dimensional problem to study the influence of the active learning parameters, namely, the number of pretrain steps N_0 , the reject threshold γ_v , and the retrain threshold γ_L . Every setting is repeated 10 times with different random seeds to account for the stochastic nature of the MCMC algorithm. The GP hyperparameter optimisation is performed 200 times with different initialisations to ensure convergence to the global minimum. The result of a grid search over these parameters is shown in Fig. 6. The Wasserstein 2-distance with respect to the reference solution from Section 5.1 is used as a metric to evaluate the accuracy of the active learning approach. While there is little variance in the results in Figs. 6(a) and 6(b), the results in Figs. 6(c) and 6(d) show the impact of the initial training points N_0 : the more training points, the smaller the approximation error. It also

becomes clear that the retrain threshold γ_L only has a minor impact on the accuracy of the inference.

The reject threshold γ_v emerges as the dominant factor when comparing the results of Fig. 6(a) - Fig. 6(d), each of which represent a different value of γ_v . No matter the number of initial MCMC steps N_0 or the retrain threshold γ_L , the accuracy of the active learning approach clearly depends on the reject threshold γ_v . These results also suggest that the active learning approach is robust with respect to the initial hyperparameter estimation, as the Wasserstein 2-distance appears to be stable across the different values of the retrain threshold. We set the active learning parameters to $N_0 = 20$, $\gamma_L = 2.5$, and $\gamma_v \in \{1.0, 5.0, 20.0\}$ for the remaining experiments.

5.3. Online vs offline learning

This section explores the training strategies for surrogate models, specifically contrasting the online active learning approach with conventional offline a priori methods. In a priori learning, the training dataset is predetermined before the inference process begins, without considering the data observed during this process. We examine three predominant strategies for compiling this dataset:

- Grid: the training data forms a structured grid;
- LHS: the training data is uniformly distributed and sampled according to a Latin hypercube design (McKay et al., 1979);
- Prior: the training data is sampled from the prior distribution of the material parameters.

Despite the simplicity of sampling from the Gaussian prior, both Grid and LHS strategies encounter an implementation challenge due to the prior's infinite support. Attempting to impose a uniform distribution across this boundless domain inherently results in an improper distribution with zero density, rendering traditional sampling methods impractical. To address this, we implement a cutoff mechanism to define a finite bounding box in the latent space, enabling meaningful sampling for Grid and LHS methods. We set this cutoff at three standard deviations from the prior mean in each dimension, prioritising comprehensive coverage over the prior's support. Although this approach may slightly compromise predictive accuracy within the denser regions of the prior, it significantly enhances model performance in instances where a notable discrepancy exists between the prior and posterior distributions. The active learning approach of Section 3, on the other hand, dynamically adjusts the training dataset based on the observed data, thereby adapting to the posterior distribution. Therefore, it self-determines the amount and location of the training data points. All four sampling strategies are depicted in Fig. 7.

We now want to compare the accuracy of the surrogate models trained with the four different training data collection strategies. To this end, we consider the range of latent space dimensionalites $d \in \{2, 3, 4, 5\}$ and the associated inverse problems. For the a priori learning approaches, we set the number of training data points to $N_{\text{train}} = 5^d$ for an even coverage in all directions. The GP-hyperparameters are estimated and fixed for the subsequent inference process. Each setting is repeated 10 times with different ground truths to account for the stochastic nature of the training process and the MCMC. The GP hyperparameter optimisation is performed 20 times with different initialisations.

Fig. 8 presents the outcomes of our comparative analysis between active learning and a priori training data collection strategies. We observe that for a latent space with dimensionality of $d = 2$ and $d = 3$, the performance across all strategies is relatively similar, albeit with grid sampling showing a slight disadvantage. This parity shifts as we move to higher dimensions ($d = 4$ and $d = 5$), where the active learning strategy outperforms the a priori methods. Moreover, a priori learning approaches demonstrate a broader variance in approximation error; some iterations are nearly as effective as active

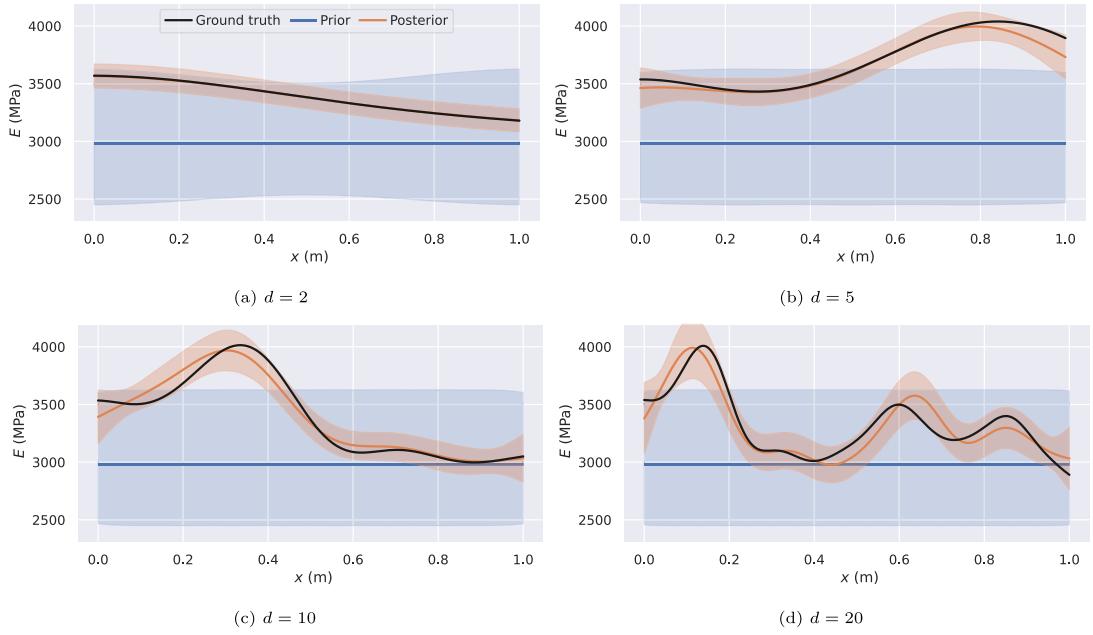


Fig. 5. Posterior realisations of the spatially varying initial stiffness E for different numbers of RBFs and respective length scales. The shaded areas represent the 95% credible intervals of their respective distribution.

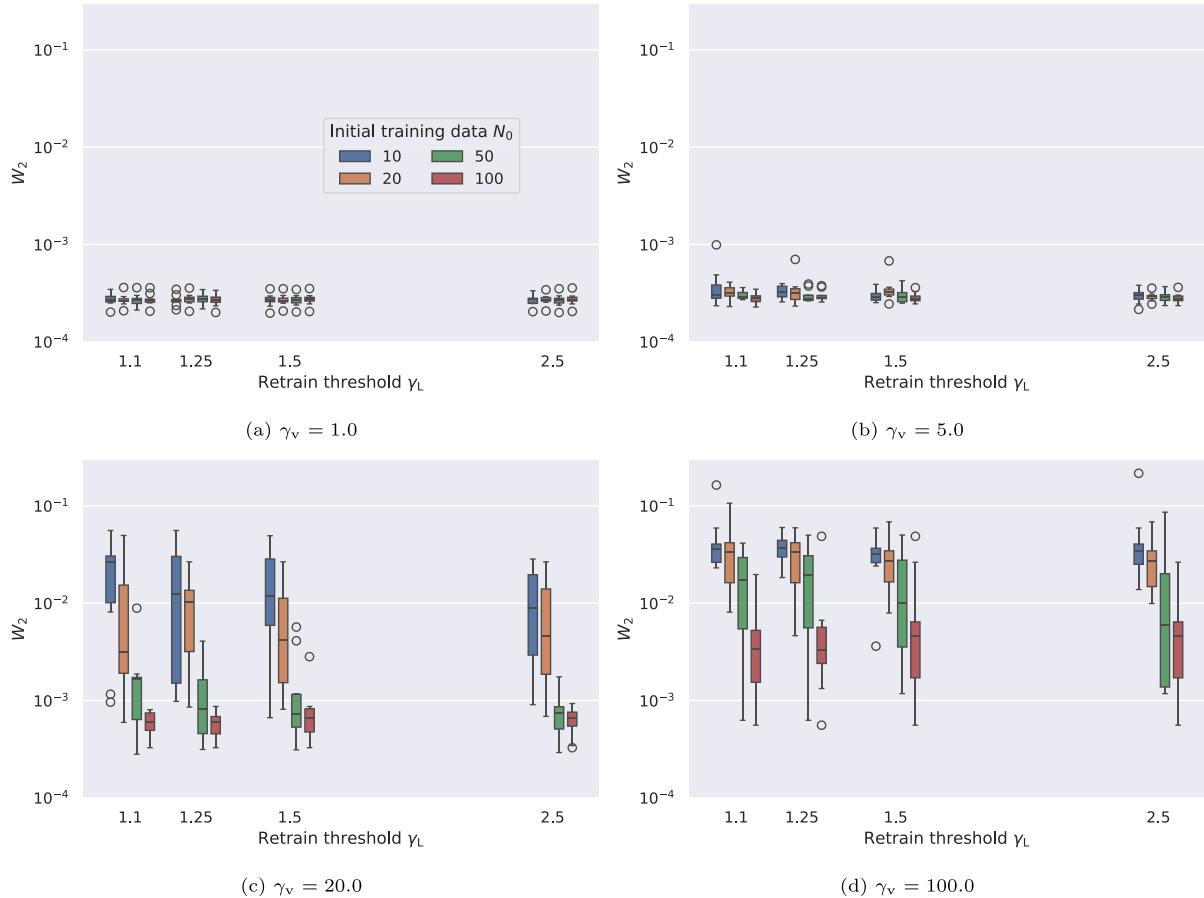


Fig. 6. Accuracy of the active learning approach on the 5 dimensional problem for different values of the reject threshold γ_v in subfigures (a)–(d). Each subfigure shows the Wasserstein 2-distance W_2 between the reference and approximate posterior distribution for different values of the retrain threshold γ_L . The number of initial training data N_0 is colour coded.

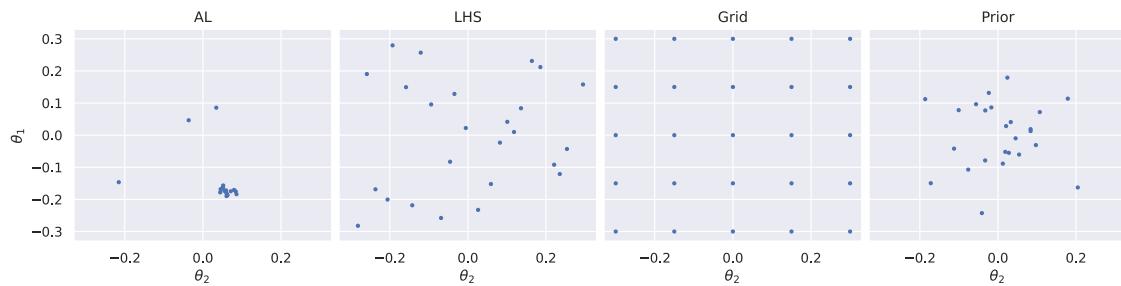


Fig. 7. Visualisation of the training data for the 2-dimensional problem. The left subplot shows the training data for the active learning approach, while the remaining subplots show the training data for the a priori learning approaches.

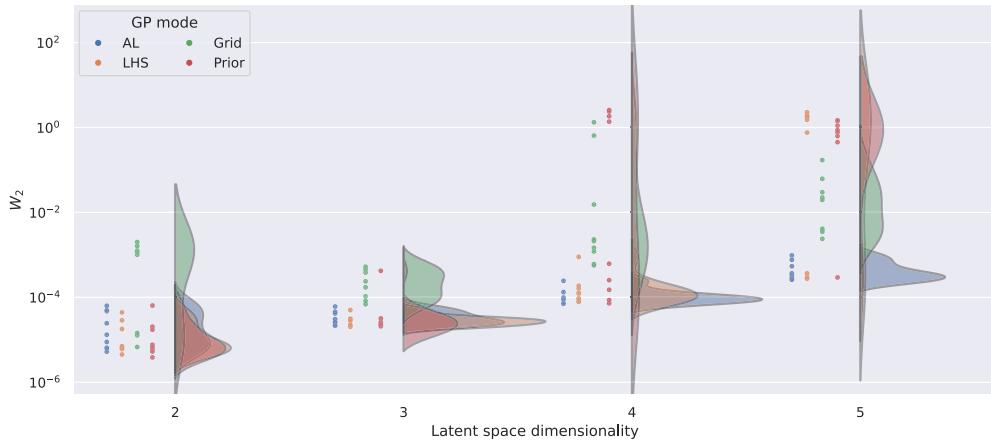


Fig. 8. Approximation error of the training data collection strategies for a range of latent space dimensionalities. The half-violins stem from a kernel density estimate of the scattered data in log-space. The offline learning approaches each get $N_{train} = 5^d$ training data, while the active learning approach gathers data according to Algorithm 1.

learning, while others fall markedly short. This variability is largely attributable to the spatial distribution of training data points with respect to regions of high posterior probability. In lower-dimensional spaces, the probability of a priori training data aligning with these critical regions remains high. However, as the dimensionality increases, the likelihood of such fortuitous alignment decreases significantly. This diminishing probability underlies the observed trends in approximation error: outcomes are either closely aligned with those of the active learning approach or significantly divergent, with few instances falling in between. These results underscore the increasing advantage of active learning in higher-dimensional spaces, attributed to its strategic adaptability in navigating the complexities introduced by the curse of dimensionality.

The trends observed with the exponential increase in number of training points can be further corroborated by setting the number of training data points to be as similar as possible for all strategies, as determined by the active learning approach. For this purpose, we set the number of training data points to $N_{train} = \lceil \sqrt{N_{train,AL}} \rceil^d$ for the a priori learning approaches, where $\lceil \cdot \rceil$ denotes smallest integer greater than the argument. This is slightly favourable for the a priori learning approaches, as $N_{train} \geq N_{train,AL}$. Nonetheless, the active learning now clearly outperforms the a priori approaches for all latent space dimensionalities. Results can be seen in Fig. 9.

As to why the a priori learning approaches perform worse than the active learning approach, it is helpful to inspect the different posteriors resulting from a single run with $d = 5$ in Fig. 10. The regions close to the boundaries are well represented by the LHS surrogate model in Fig. 10(b), whereas the inner part of the domain exposes a significant approximation error compared to the AL solution Fig. 10(a). This discrepancy is further highlighted in the pair plots of the first two latent variables in the bottom row (Figs. 10(c) and 10(d)). The posterior from

the AL surrogate model is in good agreement with the FEM posterior. The training data samples are concentrated in regions of high posterior probability, as indicated by their darker colours.

On the contrary, the LHS surrogate model fails to match the reference posterior, as illustrated by the misalignment of the red and cyan samples in Fig. 10(d). A closer look at the training data reveals the shortcoming of the LHS strategy: while the 2-dimensional projection suggest a good coverage of the latent space, the samples are not concentrated in regions of high posterior density. Only a handful of points ended up in the region of interest, leading to a poor approximation of the posterior distribution. This misalignment is naturally more pronounced in higher-dimensional settings, but can already introduce large approximation errors for a latent space dimensionality as low as $d = 3$.

5.4. Performance and accuracy comparison

After exposing the deficiencies of the a priori training data collection strategies, we settle on the active learning approach for the remainder of the experiments. In this section, we consider the MALA alongside the RWM algorithm, and compare their impact on surrogate model construction, sampling efficiency, and accuracy. We investigate the performance of the MALA and the RWM algorithm across a range of latent space dimensionalities, from $d = 2$ to $d = 15$. Every setting was repeated with 50 different ground-truth realisations (Fig. 4) to account for the inherent randomness of the MCMC algorithms. The GP hyperparameter optimisation is performed 50 times with different initialisations.

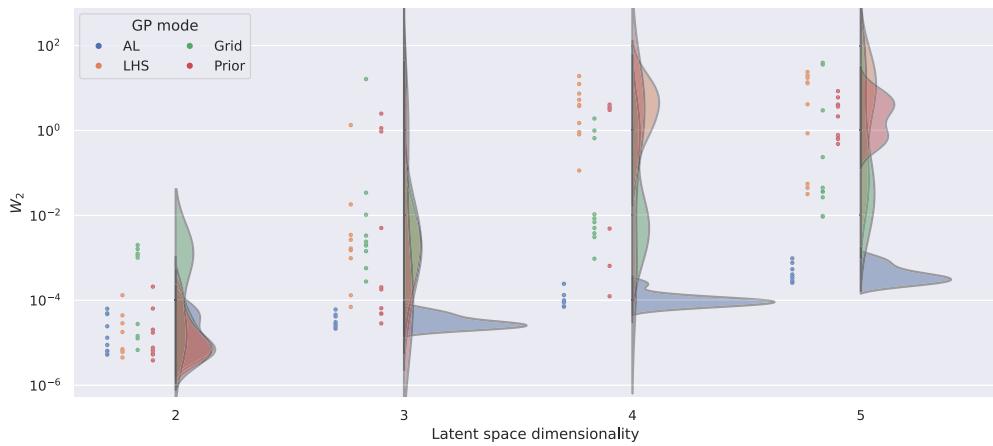


Fig. 9. Approximation error of the training data collection strategies for a range of latent space dimensionalities. The violins stem from a kernel density estimate of the scattered data in log-space. All strategies get the same amount of training data, as determined by the active learning approach.

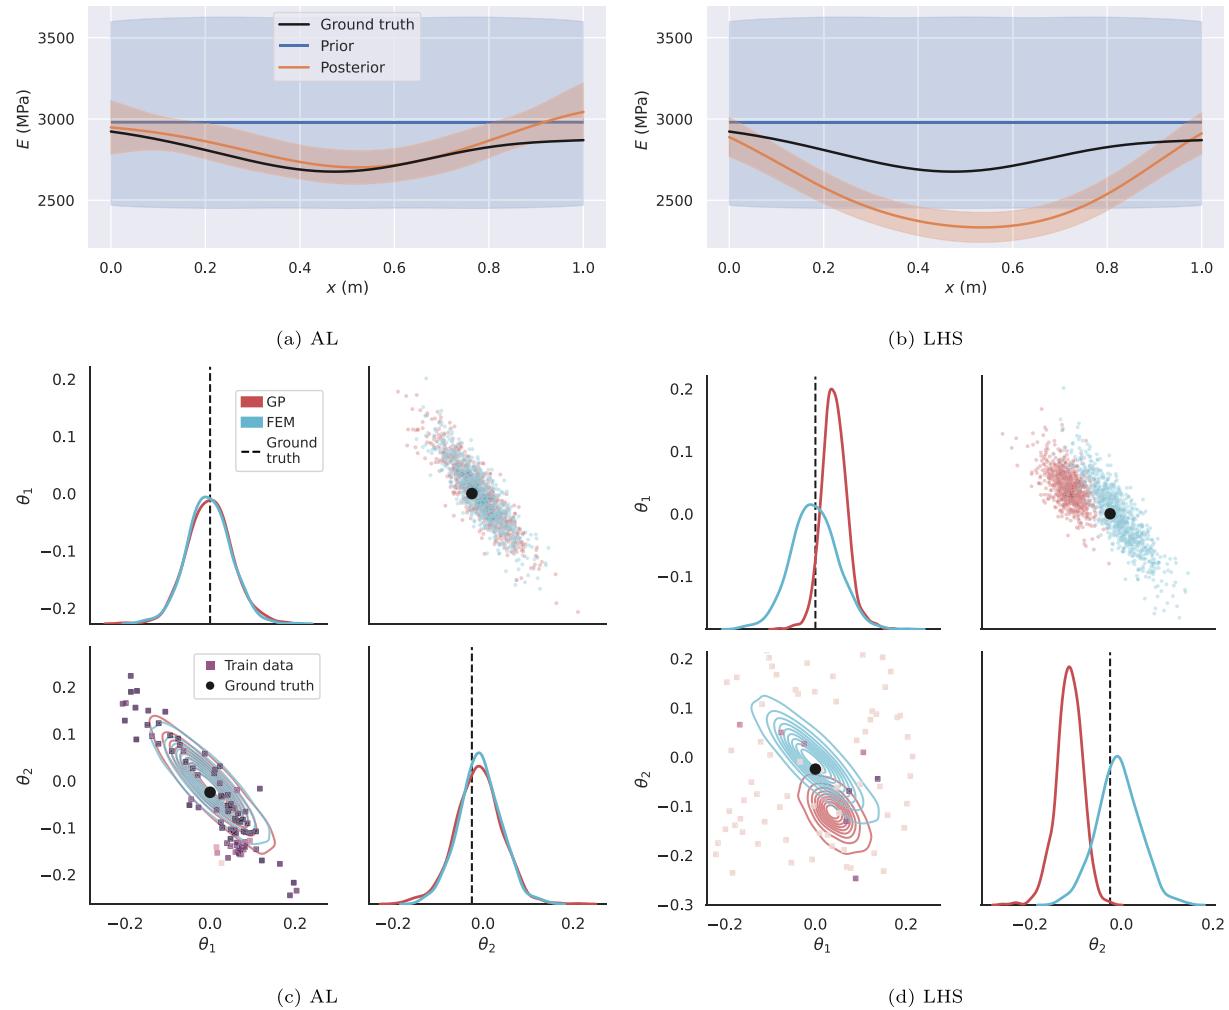


Fig. 10. Comparison of the surrogate models trained with the active learning and LHS strategies for a realisation of the problem with a 5-dimensional latent space. The top row shows the true, prior, and posterior fields, while the bottom row shows the pair plots of the first two latent space parameters. The training data are colour coded—the darker the colour, the higher the posterior density.

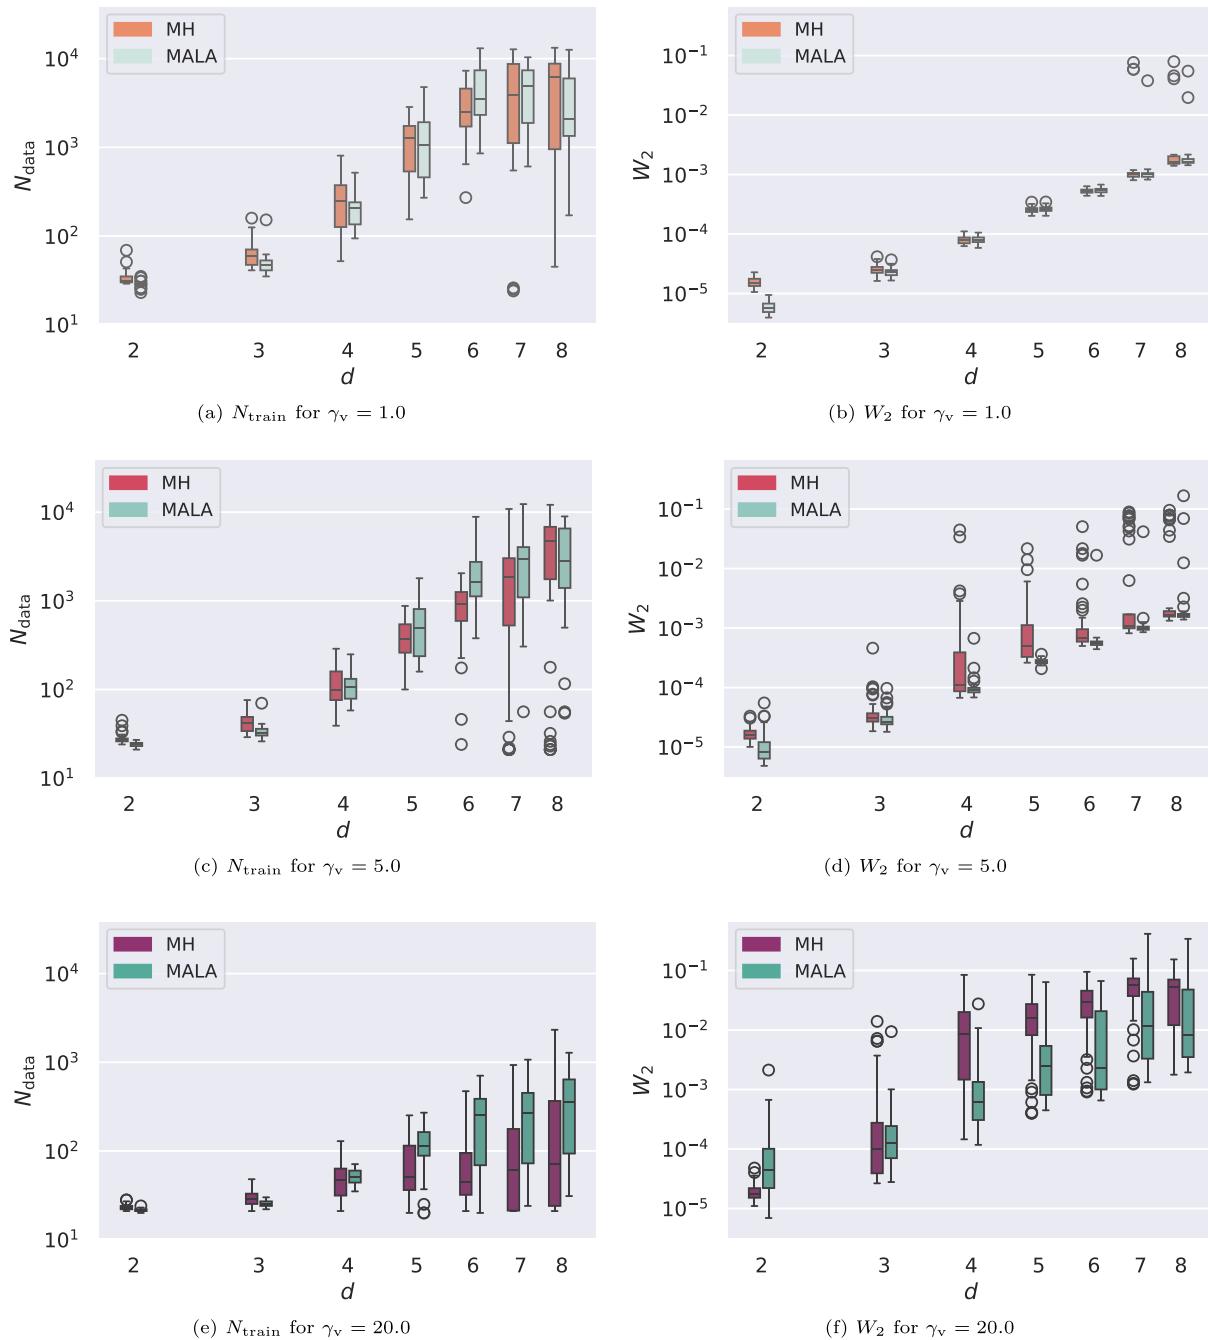


Fig. 11. Number of FEM model calls N_{train} (left column) and Wasserstein 2-distance W_2 (right column) for the RWM algorithm and the MALA across a range of latent space dimensionalities d for different reject thresholds γ_v .

5.4.1. Surrogate model construction

We begin by evaluating the number of training data points across various latent space dimensions and for three distinct reject thresholds γ_v . Since the proposal mechanism of the sampling algorithms drives the training data acquisition, the MALA and the RWM algorithm construct different surrogate models.

The MALA requires fewer training data points than the RWM algorithm across all latent space dimensionalities (Fig. 11) under stringent uncertainty requirements. With the added value of the gradient information, the MALA generates more informative proposals, which are more likely to be accepted. Therefore, the MALA requires fewer steps to generate sufficient training data for the surrogate model. The MALA advantage regarding the training data is persistent across all d ,

shown by the constant offset to RWM in log-space. The slopes of both algorithms, however, are similar, indicating that they are subject to the curse of dimensionality in a similar manner.

The right column of Fig. 11 details the accuracy of the inferences. Under strict uncertainty controls ($\gamma_v = 1.0$), both RWM and MALA achieve similar levels of accuracy in approximating the posterior distribution. Outliers below the trend line in Fig. 11(a) correspond to the outliers with high W_2 in Fig. 11(b), highlighting the tradeoff between accuracy and efficiency. These outliers suggest a superior robustness of active learning with the RWM, as they exclusively occurred with the MALA.

With looser uncertainty requirements these outliers become more prevalent (Figs. 11(c) and 11(d)) and eventually they become the norm

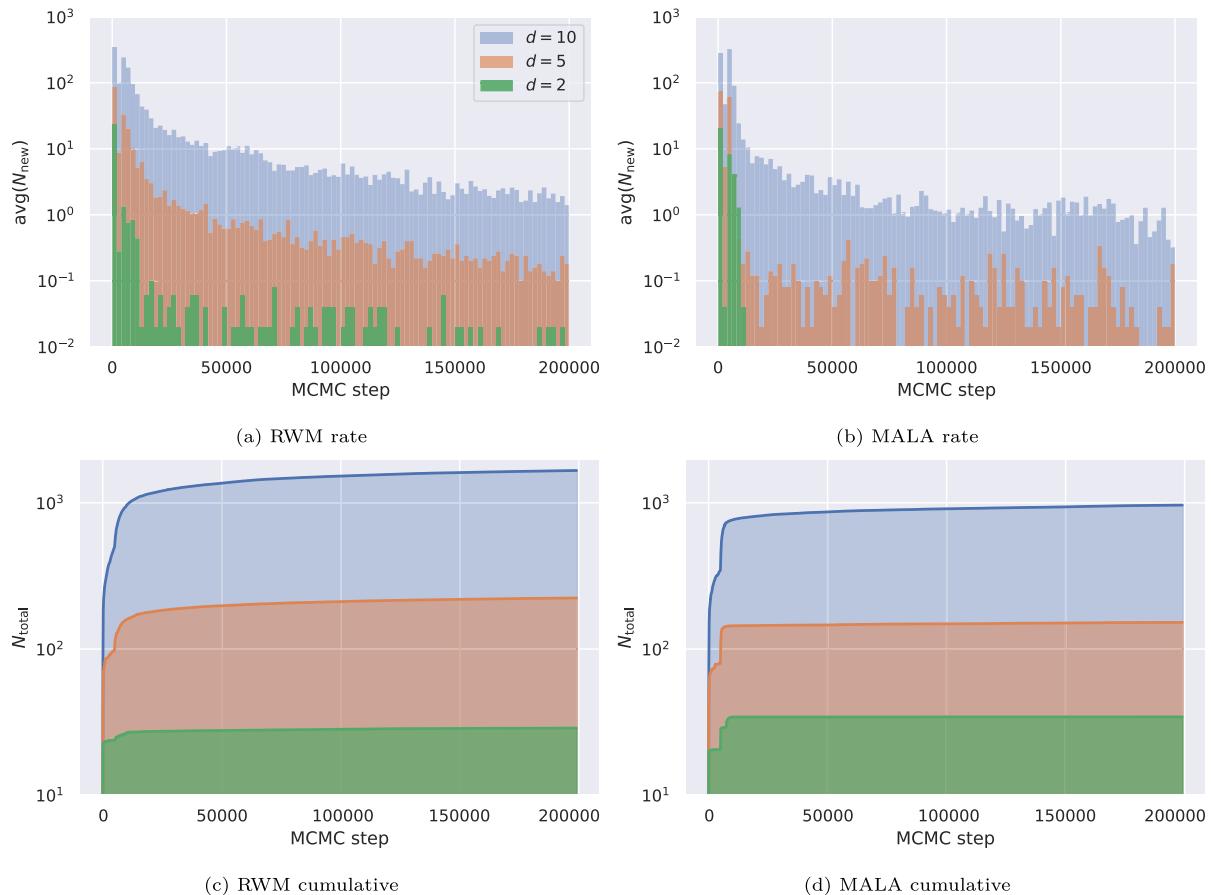


Fig. 12. Level of adaptation of the MCMC algorithms RWM in Fig. 12(a) and MALA in Fig. 12(b). Each bar represents the number of new data points collected within intervals of 2000 MCMC steps, averaged over 50 realisations. The latent space dimensionality is colour coded.

(Figs. 11(e) and 11(f)). What happens in these MALA runs is that the initial MCMC samples are too far removed from the posterior's high-density regions. These samples then poorly represent the likelihood's complexity, potentially leading to an overestimated GP length scale and an overly confident surrogate model that ceases to collect more training data. While the RWM algorithm encounters the same issue, it is less pronounced due to the more exploratory nature of the random walk proposal: eventually the algorithm moves towards the high-density regions of the posterior by chance, leading to further collection of training data. The MALA, however, is misled by an erroneous surrogate model and its poor approximation of the gradient and cannot correct itself.

To summarise the findings, the RWM algorithm and the MALA perform similarly both in terms of the number of training data points and the accuracy of the inference when strong requirements are imposed on the uncertainty in the surrogate model. When the uncertainty tolerance is relaxed, the RWM still produces accurate surrogates, underlining its superior robustness, while the accuracy of the surrogates trained with the MALA deteriorates.

5.4.2. Adaptation of the surrogate model

While the MCMC algorithms are guaranteed to converge to the distribution given by the surrogate model, the surrogate model itself must also converge to the true likelihood function. To investigate the level of adaptation of the surrogate model, we analyse the number of new data points collected throughout the MCMC chain. Fig. 12 shows the rate of data points collected per interval of 2000 MCMC steps for the

RWM (Fig. 12(a)) and MALA (Fig. 12(b)) algorithms, averaged over 50 MCMC chains each. The cumulative number of data points collected for both algorithms is shown in Fig. 12(c) and Fig. 12(d). Three conclusions can be drawn: (i) the rate of training data collection decreases in all cases, (ii) the rate is higher for the RWM algorithm than for the MALA, (iii) the rate depends on latent space dimensionality—the higher the dimensionality, the more data points are needed at all stages of the chain. Even though the rate reduces to one new data point per 2000 MCMC steps towards the end of the $d = 10$ case, this dependence suggests a stationary surrogate model becomes more difficult to obtain with increasing latent space dimensionality.

5.4.3. Sampling performance given a trained surrogate model

Finally, we compare the sampling accuracy and efficiency of the RWM algorithm and the MALA when using the same surrogate model. In this case, we are using the one constructed by the RWM algorithm with $\gamma_v = 1.0$ and prohibit further training data collection. The numerical experiment is repeated for the same 50 distinct realisations of the inverse problem per latent space dimensionality as in Section 5.4.1. There is no clear difference in accuracy of the posterior approximation between the MALA and the RWM algorithm for any given dimensionality. Fig. 13 shows the Wasserstein 2-distance across a range of latent space dimensionalities.

Starting from $d = 10$, we observe an increasing number of outliers with high W_2 values. These outliers represent chains that have completely diverged from regions with high posterior density. These chains navigated to largely unexplored regions of the parameter space,

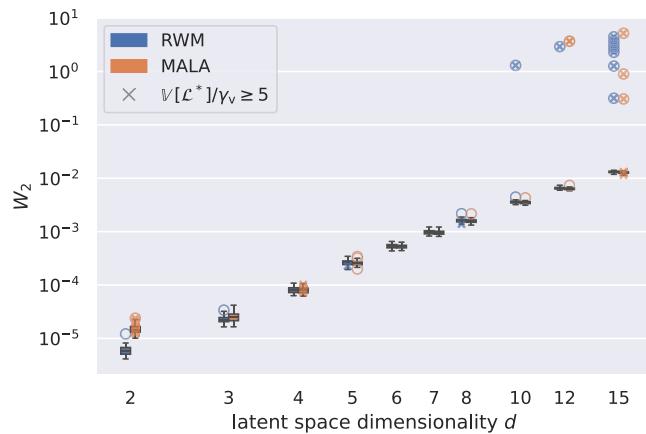


Fig. 13. Accuracy of the approximate inference when the RWM algorithm and the MALA use the same surrogate model, across a range of latent space dimensionalities. The crosses in indicate the outliers with predictive variance.

as indicated by the large predictive variance encountered during the respective MCMC run. Runs with at least one occurrence of a maximum predictive variance five times larger than γ_v are marked with a cross, which match well with the outliers in Fig. 13. This phenomenon occurs for both algorithms, but is more pronounced for the RWM algorithm. It underscores the shortcomings of a priori trained surrogate models in high-dimensional settings, as even a model tailored to a given posterior cannot guarantee the accuracy of the inference.

6. Conclusions

In this research, we have examined the integration of surrogate modelling with Markov chain Monte Carlo (MCMC) sampling techniques to improve the efficiency of Bayesian model calibration. We introduced a simple, scalable one-dimensional bar problem equipped with a non-linear constitutive model and spatially varying stiffness. The stiffness field is discretised with radial basis functions, the number of which can be adapted to control the dimensionality of the latent space. This problem serves as our test bed to explore the impacts of various methodological decisions associated with MCMC and surrogate applications.

We employed a Gaussian Process (GP) surrogate model to approximate the log likelihood function and evaluated several traditional a priori training strategies: Latin hypercube sampling, uniform grid sampling, and sampling from the prior distribution. Furthermore, we proposed an active learning strategy that adaptively selects training data points based on the MCMC path, leveraging the GP model's inherent uncertainty quantification capabilities. Our findings underscore the limitations of traditional a priori methods and demonstrate the superiority of active learning in balancing computational costs with accuracy, particularly for problems with higher stochastic dimensionality.

Comparative analysis of the random walk Metropolis (RWM) and Metropolis-adjusted Langevin algorithm (MALA) in terms of surrogate model construction, sampling efficiency, and inference accuracy revealed that both algorithms perform comparably under stringent surrogate model uncertainty requirements. The increased sampling efficiency of the MALA was beneficial for the training of the surrogate model. The higher quality of the MALA samples lead to the construction of a surrogate model with similar accuracy to the RWM algorithm, but with fewer FEM model evaluations. However, under relaxed conditions, accuracy was better maintained with the RWM algorithm than with the MALA, underscoring the robustness of the RWM algorithm. When a predefined surrogate model was used, the choice of the sampling algorithm did not affect inference accuracy. However, there is a notable

chance for both algorithms to depart from the regions of high surrogate-model certainty and accuracy. These divergent chains could not return to the high-density regions, even for surrogate models tailored to the posterior distribution, which further emphasised the shortcomings of a priori training.

This study suggest that related works (Deveney et al., 2023; Hu et al., 2024; Thomas et al., 2022; Chen et al., 2023) incorrectly identify the MCMC algorithm as a bottleneck in the inference process. The intensive training data requirements for constructing a surrogate model in high-dimensional scenarios render the MCMC-with-surrogate approach impractical before any significant benefits of advanced sampling algorithms can be realised. With a hypothetical simulator that takes 10 min to run, the compilation of the training dataset for a 15-dimensional latent space would require roughly 70 days. Any gains in sampling efficiency from advanced MCMC algorithms would be overshadowed by the prohibitive costs of surrogate model construction. This finding is at odds with common practice, where engineers often put significant resources into the integration and analysis of advanced MCMC algorithms. This study clearly suggests prioritising the surrogate model construction over the choice of sampling algorithm. Further research into multi-fidelity MCMC schemes that exploit a hierarchy of models with varying complexity (Zhang et al., 2018; Meng, 2020; Kennedy and O'Hagan, 2000; Torzoni et al., 2023) and nonlinear dimensionality reduction techniques (Dasgupta et al., 2024) could provide viable strategies to mitigate the curse of dimensionality.

CRedit authorship contribution statement

Leon Riccius: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Conceptualization. **Iuri B.C.M. Rocha:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Joris Bierkens:** Writing – review & editing, Methodology. **Hanne Kekkonen:** Writing – review & editing, Methodology. **Frans P. van der Meer:** Writing – review & editing, Supervision, Conceptualization.

Code availability

The code used in this study is available at github.com/SLIMM-Lab/mcmc-with-surrogates.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work is supported by the TU Delft AI Labs programme through the SLIMM lab.

Data availability

Data will be made available on request.

References

Bierkens, J., Fearnhead, P., Roberts, G., 2019. The Zig-Zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Statist.* 47, 1288–1320. <http://dx.doi.org/10.1214/18-AOS1715>, URL: <https://projecteuclid.org/journals/annals-of-statistics/volume-47/issue-3/The-Zig-Zag-process-and-super-efficient-sampling-for-Bayesian/10.1214/18-AOS1715.full>.

Chen, Z., Jin, P., Li, R., Qi, Y., Cai, G., 2023. Parameter identification of elastoplastic model for CuCrZr alloy by the neural network-aided Bayesian inference. *Fatigue Fract. Eng. Mater. Struct.* 46, 2319–2337. <http://dx.doi.org/10.1111/ffe.14000>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/ffe.14000>.

Ching, J., Chen, Y.-C., 2007. Transitional Markov chain Monte Carlo method for Bayesian model updating, model class selection, and model averaging. *J. Eng. Mech.* 133, 816–832.

Chong, A., Lam, K.P., 2017. A comparison of MCMC algorithms for the Bayesian calibration of building energy models for building simulation 2017 conference. In: Building Simulation Conference Proceedings. vol. 2, pp. 582–591. <http://dx.doi.org/10.26868/25222708.2017.336>.

Dasgupta, A., Patel, D.V., Ray, D., Johnson, E.A., Oberai, A.A., 2024. A dimension-reduced variational approach for solving physics-based inverse problems using generative adversarial network priors and normalizing flows. *Comput. Methods Appl. Mech. Engrg.* 420, 116682, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0045782523008058>.

Deveney, T., Mueller, E.H., Shardlow, T., 2023. Deep surrogate accelerated delayed-acceptance Hamiltonian Monte Carlo for Bayesian inference of spatio-temporal heat fluxes in rotating disc systems. *SIAM/ASA J. Uncertain. Quantif.* 11, 970–995, URL: <https://pubs.siam.org/doi/10.1137/22M1513113>.

Dinkel, M., Geitner, C.M., Rei, G.R., Nitzler, J., Wall, W.A., 2024. Solving Bayesian inverse problems with expensive likelihoods using constrained Gaussian processes and active learning. *Inverse Problems* 40, 095008. <http://dx.doi.org/10.1088/1361-6420/ad5eb4>, URL: <https://iopscience.iop.org/article/10.1088/1361-6420/ad5eb4>.

Dobrilla, S., Lunardelli, M., Nikolić, M., Lowke, D., Rosić, B., 2023. Bayesian inference of mesoscale mechanical properties of mortar using experimental data from a double shear test. *Comput. Methods Appl. Mech. Engrg.* 409, 115964, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0045782523008077>.

Drovandi, C.C., Moores, M.T., Boys, R.J., 2018. Accelerating pseudo-marginal MCMC using Gaussian processes. *Comput. Statist. Data Anal.* 118, 1–17.

Duane, S., Kennedy, A., Pendleton, B.J., Roweth, D., 1987. Hybrid Monte Carlo. *Phys. Lett. B* 195, 216–222, URL: <https://linkinghub.elsevier.com/retrieve/pii/037026938791197X>.

Fletcher, R., 2000. Practical methods of optimization. *Pr. Methods Optim.* URL: <https://onlinelibrary.wiley.com/doi/book/10.1002/9781118723203>.

Gelman, A., Gilks, W.R., Roberts, G.O., 1997. Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Probab.* 7, 110–120, URL: <https://projecteuclid.org/journals/annals-of-applied-probability/volume-7/issue-1/Weak-convergence-and-optimal-scaling-of-random-walk-Metropolis-algorithms/10.1214/aoap/1034625254.full>.

Ghanem, R.G., Spanos, P.D., 1991. Stochastic Finite Elements: A Spectral Approach. Springer New York, New York, NY, URL: <http://link.springer.com/10.1007/978-1-4612-3094-6>.

Girolami, M., Calderhead, B., 2011. Riemann manifold langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73, 123–214, URL: www.ucl.ac.uk/statistics/.

Goodman, J., Weare, J., 2010. Ensemble samplers with affine invariance. *Commun. Appl. Math. Comput. Sci.* 5, 65–80.

Green, P.J., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82, 711, URL: <https://www.jstor.org/stable/2337340?origin=crossref>.

Hou, D., Hassan, I., Wang, L., 2021. Review on building energy model calibration by Bayesian inference. *Renew. Sustain. Energy Rev.* 143, 110930, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032121002239>.

Hu, Y., Abuseda, M., Algheli, A., Holdheim, S., Fisher, T.S., 2024. Surrogate-accelerated Bayesian framework for high-temperature thermal diffusivity characterization. *Comput. Methods Appl. Mech. Engrg.* 418, 116459. <http://dx.doi.org/10.1016/j.cma.2023.116459>.

Kandasamy, K., Schneider, J., Póczos, B., 2015. Bayesian active learning for posterior estimation. In: IJCAI International Joint Conference on Artificial Intelligence. vol. 2015-Janua, pp. 3605–3611.

Kennedy, M.C., O'Hagan, A., 2000. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* 87, 1–13.

Kennedy, M.C., O'Hagan, A., 2001. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 63, 425–464, URL: <https://onlinelibrary.wiley.com/doi/full/10.1111/1467-9868.00294>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.00294>.

Kiran, R.P., Das, A., Bansal, S., 2025. A state-of-the-art review of Bayesian finite element model updating techniques for structural systems. *Probabilistic Eng. Mech.* 80, 103761.

Li, C.-C., Kiureghian, A.D., 1993. Optimal discretization of random fields. *J. Eng. Mech.* 119, 1136–1154. [http://dx.doi.org/10.1061/\(ASCE\)0733-9399\(1993\)119:6\(1136\)](http://dx.doi.org/10.1061/(ASCE)0733-9399(1993)119:6(1136)), <https://ascelibrary.org/doi/abs/10.1061/%28ASCE%290733-9399%281993%29119%3A6%281136%29>.

Lye, A., Cicirello, A., Patelli, E., 2022. An efficient and robust sampler for Bayesian inference: Transitional ensemble Markov chain Monte Carlo. *Mech. Syst. Signal Process.* 167, 108471.

Mackay, D.J.C., 1998. Introduction to Gaussian processes. In: Neural Networks and Machine Learning. Barber and Williams, URL: <http://www.cs.toronto.edu/~radford/>.

Marsili, F., Croce, P., Friedman, N., Formichi, P., Landi, F., 2017. Seismic reliability assessment of a concrete water tank based on the Bayesian updating of the finite element model. *ASCE-ASME J. Risk Uncertain. Eng. Syst. Part B: Mech. Eng.* 3, URL: <https://asmedigitalcollection.asme.org/risk/article/3/2/021004/369959/Seismic-Reliability-Assessment-of-a-Concrete-Water>.

Marzouk, Y.M., Najm, H.N., 2009. Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *J. Comput. Phys.* 228, 1862–1902, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0021999108006062>.

McKay, M.D., Beckman, R.J., Conover, W.J., 1979. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 21, 239, URL: <https://www.jstor.org/stable/1268522?origin=crossref>.

Melro, A.R., Camanho, P.P., Pires, F.M.A., Pinho, S.T., 2013. Micromechanical analysis of polymer composites reinforced by unidirectional fibres: Part I-Constitutive modelling. *Int. J. Solids Struct.* 50, 1897–1905.

Meng, X., 2020. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *Karniadakis / J. Comput. Phys.* 401, 109020, URL: www.elsevier.com/locate/jcp.

Mengersen, K.L., Tweedie, R.L., 1996. Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist.* 24, 101–121, URL: <https://projecteuclid.org/journals/annals-of-statistics/volume-24/issue-1/Rates-of-convergence-of-the-Hastings-and-Metropolis-algorithms/10.1214/aos/103306201.full>.

Nouy, A., Maître, O.P.L., 2009. Generalized spectral decomposition for stochastic nonlinear problems. *J. Comput. Phys.* 228, 202–235, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0021999108004737>.

Rappel, H., Beex, L.A., Hale, J.S., Noels, L., Bordas, S.P., 2020. A tutorial on Bayesian inference to identify material parameters in solid mechanics. *Arch. Comput. Methods Eng.* 27, 361–385.

Rappel, H., Beex, L.A., Noels, L., Bordas, S.P., 2019. Identifying elastoplastic parameters with Bayes' theorem considering output error, input error and model uncertainty. *Probabilistic Eng. Mech.* 55, 28–41.

Rasmussen, C.E., Williams, C.K.I., 2006. Gaussian Processes for Machine Learning. MIT Press, pp. I–XVIII, 1–248, URL: www.GaussianProcess.org/gpm.

Roberts, G.O., Rosenthal, J.S., 1998. Optimal scaling of discrete approximations to langevin diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 60, 255–268, URL: <https://academic.oup.com/rssb/article/60/1/255/7083121>.

Roberts, G.O., Rosenthal, J.S., 2001. Optimal scaling for various Metropolis-Hastings algorithms. 16, (ISSN: 0883-4237) pp. 351–367, URL: <https://projecteuclid.org/journals/statistical-science/volume-16/issue-4/Optimal-scaling-for-various-Metropolis-Hastings-algorithms/10.1214/ss/1015346320.full>, <https://projecteuclid.org/journals/statistical-science/volume-16/issue-4/Optimal-scaling-for-various-Metropolis-Hastings-algorithms/10.1214/ss/1015346320.short>.

Roberts, G.O., Rosenthal, J.S., 2007. Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *J. Appl. Probab.* 44, 458–475, URL: https://www.cambridge.org/core/product/identifier/S0021900200117954/type/journal_article.

Roberts, G.O., Stramer, O., 2002. Langevin diffusions and Metropolis-Hastings algorithms. *Methodol. Comput. Appl. Probab.* 2002 4:4 4, 337–357, URL: <https://link.springer.com/article/10.1023/A:1023562417138>.

Rocha, I.B., Kerfriden, P., van der Meer, F.P., 2021. On-the-fly construction of surrogate constitutive models for concurrent multiscale mechanical analysis through probabilistic machine learning. *J. Comput. Phys.: X* 9, 100083.

Römer, U., Hartmann, S., Tröger, J.-A., Anton, D., Wessels, H., Flaschel, M., Lorenzis, L.D., 2025. Reduced and all-at-once approaches for model calibration and discovery in computational solid mechanics. *Appl. Mech. Rev.* 77, URL: <https://asmedigitalcollection.asme.org/appliedmechanicsreviews/article/77/4/040801/1201974/Reduced-and-All-At-Once-Approaches-for-Model>.

Simoen, E., Papadimitriou, C., Lombaert, G., 2013. On prediction error correlation in Bayesian model updating. *J. Sound Vib.* 332, 4136–4152, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0022460X13002514>.

Straub, D., Papaioannou, I., 2015. Bayesian updating with structural reliability methods. *J. Eng. Mech.* 141, 04014134.

Stuart, A.M., 2010. Inverse problems: A Bayesian perspective. *Acta Numer.* 19, 451–459, URL: <https://doi.org/10.1017/S0962492910000061>.

Sudret, B., Kiureghian, A.D., 2000. Stochastic finite element methods and reliability: a state-of-the-art report. URL: <https://books.google.nl/books?id=PNNztgAACAAJ>.

Thomas, A.J., Barocio, E., Bilionis, I., Pipes, R.B., 2022. Bayesian inference of fiber orientation and polymer properties in short fiber-reinforced polymer composites. *Compos. Sci. Technol.* 228, 109630, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0266353822003724>.

Torzoni, M., Manzoni, A., Mariani, S., 2023. A multi-fidelity surrogate model for structural health monitoring exploiting model order reduction and artificial neural networks. *Mech. Syst. Signal Process.* 197, 110376, URL: <http://creativecommons.org/licenses/by/4.0/>.

Uribe, F., Papaioannou, I., Betz, W., Straub, D., 2020. Bayesian inference of random fields represented with the Karhunen–Loëve expansion. *Comput. Methods Appl. Mech. Engrg.* 358, 112632, URL: <https://linkinghub.elsevier.com/retrieve/pii/S004578251930516X>.

del Val, A., Maître, O.P.L., Magin, T.E., Chazot, O., Congedo, P.M., 2022. A surrogate-based optimal likelihood function for the Bayesian calibration of catalytic recombination in atmospheric entry protection materials. *Appl. Math. Model.* 101, 791–810.

Vigliotti, A., Csányi, G., Deshpande, V.S., 2018. Bayesian inference of the spatial distributions of material properties. *J. Mech. Phys. Solids* 118, 74–97.

Wu, L., Zulueta, K., Major, Z., Arriaga, A., Noels, L., 2020. Bayesian inference of nonlinear multiscale model parameters accelerated by a deep neural network. *Comput. Methods Appl. Mech. Engrg.* 360, URL: www.sciencedirect.comwww.elsevier.com/locate/cma.

Zhang, J., Man, J., Lin, G., Wu, L., Zeng, L., 2018. Inverse modeling of hydrologic systems with adaptive multifidelity Markov chain Monte Carlo simulations. *Water Resour. Res.* 54, 4867–4886, URL: <https://onlinelibrary.wiley.com/doi/full/10.1029/2018WR022658>, <https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2018WR022658>.